

A DEIM Induced CUR Factorization ^{*}

D. C. SORENSEN[†] AND M. EMBREE[‡]

e-mail: sorensen@rice.edu, embree@vt.edu

September 22, 2015

Abstract

We derive a CUR approximate matrix factorization based on the Discrete Empirical Interpolation Method (DEIM). For a given matrix \mathbf{A} , such a factorization provides a low rank approximate decomposition of the form $\mathbf{A} \approx \mathbf{CUR}$, where \mathbf{C} and \mathbf{R} are subsets of the columns and rows of \mathbf{A} , and \mathbf{U} is constructed to make \mathbf{CUR} a good approximation. Given a low-rank singular value decomposition $\mathbf{A} \approx \mathbf{VSW}^T$, the DEIM procedure uses \mathbf{V} and \mathbf{W} to select the columns and rows of \mathbf{A} that form \mathbf{C} and \mathbf{R} . Through an error analysis applicable to a general class of CUR factorizations, we show that the accuracy tracks the optimal approximation error within a factor that depends on the conditioning of submatrices of \mathbf{V} and \mathbf{W} . For very large problems, \mathbf{V} and \mathbf{W} can be approximated well using an incremental QR algorithm that makes only one pass through \mathbf{A} . Numerical examples illustrate the favorable performance of the DEIM-CUR method compared to CUR approximations based on leverage scores.

1 Introduction

This work presents a new CUR matrix factorization based upon the Discrete Empirical Interpolation Method (DEIM). A CUR factorization is a low rank approximation of a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ of the form $\mathbf{A} \approx \mathbf{CUR}$, where $\mathbf{C} = \mathbf{A}(:, \mathbf{q}) \in \mathbb{R}^{m \times k}$ is a subset of the columns of \mathbf{A} and $\mathbf{R} = \mathbf{A}(\mathbf{p}, :) \in \mathbb{R}^{k \times n}$ is a subset of the rows of \mathbf{A} . (We generally assume $m \geq n$ throughout.) The $k \times k$ matrix \mathbf{U} is constructed to assure that \mathbf{CUR} is a good approximation to \mathbf{A} . Assuming the best rank- k singular value decomposition (SVD) $\mathbf{A} \approx \mathbf{VSW}^T$ is available, the algorithm uses the DEIM index selection procedure, $\mathbf{q} = \text{DEIM}(\mathbf{V})$ and $\mathbf{p} = \text{DEIM}(\mathbf{W})$, to determine \mathbf{C} and \mathbf{R} . The resulting approximate factorization is nearly as accurate as the best rank- k SVD, with

$$\|\mathbf{A} - \mathbf{CUR}\| \leq (\eta_p + \eta_q) \sigma_{k+1},$$

where σ_{k+1} is the first neglected singular value of \mathbf{A} , $\eta_p \equiv \|\mathbf{V}(\mathbf{p}, :)^{-1}\|$, and $\eta_q \equiv \|\mathbf{W}(\mathbf{q}, :)^{-1}\|$.

Here and throughout, $\|\cdot\|$ denotes the vector 2-norm and the matrix norm it induces, and $\|\cdot\|_F$ is the Frobenius norm. We use MATLAB notation to index vectors and matrices, so that, e.g., $\mathbf{A}(\mathbf{p}, :)$ denotes the k rows of \mathbf{A}

^{*}This work was supported in part by AFOSR grant FA9550-12-1-0155 and by NSF grant CCF-1320866.

[†]Department of Computational and Applied Mathematics, MS 134, Rice University, Houston, Texas 77005-1892.

[‡]Department of Mathematics, 225 Stanger Street 0123, Virginia Tech, Blacksburg, Virginia 24061

whose indices are specified by the entries of the vector $\mathbf{p} \in \mathbb{N}^k$, while $\mathbf{A}(:, \mathbf{q})$ denotes the k columns of \mathbf{A} indexed by $\mathbf{q} \in \mathbb{N}^k$.

The CUR factorization is an important tool for handling large-scale data sets, offering two advantages over the SVD: when \mathbf{A} is sparse, so too are \mathbf{C} and \mathbf{R} , unlike the matrices \mathbf{V} and \mathbf{W} of singular vectors; and the columns and rows that comprise \mathbf{C} and \mathbf{R} are representative of the data (e.g., sparse, nonnegative, integer valued, etc.). The following simple example, adapted from Mahoney and Drineas [22, Fig. 1b], illustrates the latter advantage. Construct $\mathbf{A} \in \mathbb{R}^{2 \times n}$ so that its first $n/2$ columns have the form

$$\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

and the remaining $n/2$ columns have the form

$$\frac{\sqrt{2}}{2} \begin{bmatrix} -1 & 1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix},$$

where in both cases $x_1 \sim N(0, 1)$ and $x_2 \sim N(0, 4^2)$ are independent samples of normal random variables, i.e., the columns of \mathbf{A} are drawn from two different multivariate normal distributions. Figure 1 shows that the two left singular vectors, though orthogonal by construction, fail to represent the true nature of the data; in contrast, the first two columns selected by the DEIM-CUR procedure give a much better overall representation. While trivial in this two-dimensional case, one can imagine the utility of such approximations for high-dimensional data. We shall illustrate the advantages of CUR approximations with further computational examples in Section 6.

CUR-type factorizations originated with “pseudoskeleton” approximations [14] and pivoted, truncated QR decompositions [23]; in recent years many new algorithms have been proposed in the numerical linear algebra and theoretical computer science literatures. Some approaches seek to maximize the *volume* of the decomposition [14,




Figure 1: Comparison of singular vectors (left, scaled, in red) and DEIM-CUR columns (right, in blue) for a data set drawn from two multivariate normal distributions having different principal axes.

25]. Numerous other algorithms instead use *leverage scores* [5, 10, 22, 28]. These methods typically first compute a singular value decomposition¹ $\mathbf{A} = \mathbf{V}\mathbf{S}\mathbf{W}^T$ (or an approximation to it), with $\mathbf{V} \in \mathbb{R}^{m \times n}$, $\mathbf{W} \in \mathbb{R}^{n \times n}$. The leverage score for the j th row (k th column) of \mathbf{A} is the squared two-norm of the j th row of \mathbf{V} (k th row of \mathbf{W}). When scaled by the number of singular vectors, these leverage scores give probability distributions for randomly sampling the columns and rows to form \mathbf{C} and \mathbf{R} . This approach leads to probabilistic bounds on $\|\mathbf{A} - \mathbf{CUR}\|_F$ [10, 22]. In cases where \mathbf{A} has small singular values (precisely the case where one would seek a low-rank factorization), the singular vectors can be sensitive to perturbations to \mathbf{A} , making the leverage scores unstable [18]. Thus leverage scores are often computed using only the leading few singular vectors, but the choice of how many vectors to keep can be somewhat ad hoc.

The algorithm described in Sections 2 and 3 is entirely deterministic and involves few (if any) parameters. The method is supported by an error analysis in Section 4 that also applies to a broad class of CUR factorizations. This section includes an improved bound on the error constants η_p and η_q for DEIM row and column selection, which also applies to the analysis of DEIM-based model order reduction [6]. In Section 5 we propose a novel incremental QR algorithm for approximating the SVD (and potentially also approximating leverage scores). Section 6 illustrates the performance of this new CUR factorization on several examples.

In many applications one cares primarily about key columns or rows of \mathbf{A} , rather than an explicit $\mathbf{A} = \mathbf{CUR}$ factorization. The DEIM technique, which identifies rows and columns of \mathbf{A} independently, can easily be used to select only columns or rows, leading to an “interpolatory decomposition” of the form $\mathbf{A} = \mathbf{C}\hat{\mathbf{U}}$ or $\mathbf{A} = \hat{\mathbf{U}}\mathbf{R}$; such factorizations have the advantage that $\hat{\mathbf{U}}$ can be much better conditioned than the \mathbf{U} matrix in the CUR factorization. For further details about general interpolatory decompositions, see [7, §1].

2 CUR Factorization

We are concerned with large matrices $\mathbf{A} \in \mathbb{R}^{m \times n}$ that represent nearly low-rank data, which can therefore be expressed as

$$\mathbf{A} = \mathbf{CUR} + \mathbf{F}, \quad (2.1)$$

with $\|\mathbf{F}\|$ small relative to $\|\mathbf{A}\|$. The matrix $\mathbf{C} \in \mathbb{R}^{m \times k}$ is formed by extracting k columns from \mathbf{A} , and $\mathbf{R} \in \mathbb{R}^{k \times n}$ from k rows of \mathbf{A} . The selected row and column indices are stored in the vectors $\mathbf{p}, \mathbf{q} \in \mathbb{N}^k$, so that $\mathbf{C} = \mathbf{A}(:, \mathbf{q})$ and $\mathbf{R} = \mathbf{A}(\mathbf{p}, :)$. Our choice for \mathbf{p} and \mathbf{q} is guided by knowledge of the rank- k SVD (or an approximation to it). Before detailing the method for selecting these indices, we discuss how, given \mathbf{p} and \mathbf{q} , one should construct \mathbf{U} so that \mathbf{CUR} satisfies desirable approximation properties.

As motivation, suppose for the moment that \mathbf{A} has exact rank k , and \mathbf{C} and \mathbf{R} are full-rank subsets of the columns and rows of \mathbf{A} . Now let $\mathbf{Y} \in \mathbb{R}^{m \times k}$ and $\mathbf{Z} \in \mathbb{R}^{n \times k}$ be any matrices that satisfy $\mathbf{Y}^T \mathbf{C} = \mathbf{RZ} = \mathbf{I} \in \mathbb{R}^{k \times k}$. Then \mathbf{CY}^T is a projector onto $\text{Ran}(\mathbf{C}) = \text{Ran}(\mathbf{A})$ and $(\mathbf{ZR})^T$ is a projector onto $\text{Ran}(\mathbf{R}^T) = \text{Ran}(\mathbf{A}^T)$, where $\text{Ran}(\cdot)$ denotes the range (column space). It follows that $\mathbf{CY}^T \mathbf{A} = \mathbf{A}$ and $(\mathbf{ZR})^T \mathbf{A}^T = \mathbf{A}^T$. Putting $\mathbf{U} \equiv \mathbf{Y}^T \mathbf{AZ}$ gives

$$\mathbf{CUR} = \mathbf{CY}^T \mathbf{AZR} = \mathbf{AZR} = \mathbf{A}.$$

Thus, any choice of \mathbf{Y} and \mathbf{Z} that satisfies $\mathbf{Y}^T \mathbf{C} = \mathbf{RZ} = \mathbf{I}$ gives a \mathbf{U} such that \mathbf{CUR} exactly recovers \mathbf{A} . In general different choices for \mathbf{Y} and \mathbf{Z} give different $\mathbf{U} = \mathbf{Y}^T \mathbf{AZ}$.

¹We use the nonstandard notation $\mathbf{V}\mathbf{S}\mathbf{W}^T$ for the SVD to avoid conflicts with \mathbf{U} in the standard CUR notation.

Now consider the general case (2.1). Once \mathbf{p} , \mathbf{q} , \mathbf{Y} , and \mathbf{Z} have been specified, then

$$\mathbf{U} = \mathbf{Y}^T \mathbf{A} \mathbf{Z} \quad \text{and} \quad \mathbf{F} \equiv \mathbf{A} - \mathbf{C} \mathbf{U} \mathbf{R}.$$

One might design \mathbf{Y} and \mathbf{Z} so that $\mathbf{C} \mathbf{U} \mathbf{R}$ matches the selected columns $\mathbf{C} = \mathbf{A}(:, \mathbf{q})$ and rows $\mathbf{R} = \mathbf{A}(\mathbf{p}, :)$ of \mathbf{A} exactly. This can be accomplished with *interpolatory projectors*, which we discuss in detail in the next section. For now, let $\mathbf{P} = \mathbf{I}(:, \mathbf{p}) \in \mathbb{R}^{m \times k}$ and $\mathbf{Q} = \mathbf{I}(:, \mathbf{q}) \in \mathbb{R}^{n \times k}$ be submatrices of the identity, so that $\mathbf{P}^T \mathbf{a} = \mathbf{a}(\mathbf{p})$ and $\mathbf{b}^T \mathbf{Q} = \mathbf{b}(\mathbf{q})^T$ for arbitrary vectors \mathbf{a} and \mathbf{b} of appropriate dimensions. Now define $\mathbf{Y}^T = (\mathbf{P}^T \mathbf{C})^{-1} \mathbf{P}^T$ and $\mathbf{Z} = \mathbf{Q}(\mathbf{R} \mathbf{Q})^{-1}$ (presuming $\mathbf{P}^T \mathbf{C}$ and $\mathbf{R} \mathbf{Q}$ are invertible). Then since $\mathbf{C} = \mathbf{A}(:, \mathbf{q})$ and $\mathbf{R} = \mathbf{A}(\mathbf{p}, :)$,

$$\mathbf{P}^T \mathbf{C} = \mathbf{C}(\mathbf{p}, :) = \mathbf{A}(\mathbf{p}, \mathbf{q}) \quad \text{and} \quad \mathbf{R} \mathbf{Q} = \mathbf{R}(:, \mathbf{q}) = \mathbf{A}(\mathbf{p}, \mathbf{q}),$$

so

$$\mathbf{U} = \mathbf{Y}^T \mathbf{A} \mathbf{Z} = (\mathbf{P}^T \mathbf{C})^{-1} \mathbf{P}^T \mathbf{A} \mathbf{Q} (\mathbf{R} \mathbf{Q})^{-1} = \mathbf{A}(\mathbf{p}, \mathbf{q})^{-1} \mathbf{A}(\mathbf{p}, \mathbf{q}) \mathbf{A}(\mathbf{p}, \mathbf{q})^{-1} = \mathbf{A}(\mathbf{p}, \mathbf{q})^{-1}.$$

This CUR approximation matches the \mathbf{q} columns and \mathbf{p} rows of \mathbf{A} ,

$$\mathbf{A}(:, \mathbf{q}) = \mathbf{C} \mathbf{U} \mathbf{R}(:, \mathbf{q}) \quad \text{and} \quad \mathbf{A}(\mathbf{p}, :) = \mathbf{C}(\mathbf{p}, :) \mathbf{U} \mathbf{R},$$

and, in our experiments, usually delivers a very good approximation. However, a CUR factorization with better theoretical approximation properties results from orthogonal projection, as originally suggested by Stewart [23, p. 320]; see also, e.g., Mahoney and Drineas [22]. Given a selection of indices \mathbf{p} and \mathbf{q} , again put

$$\mathbf{C} = \mathbf{A}(:, \mathbf{q}) \quad \text{and} \quad \mathbf{R} = \mathbf{A}(\mathbf{p}, :).$$

Assume that \mathbf{C} and \mathbf{R} both have full rank k , and now let $\mathbf{Y}^T = \mathbf{C}^I \equiv (\mathbf{C}^T \mathbf{C})^{-1} \mathbf{C}^T$ and $\mathbf{Z} = \mathbf{R}^I \equiv \mathbf{R}^T (\mathbf{R} \mathbf{R}^T)^{-1}$ denote left and right inverses of \mathbf{C} and \mathbf{R} . These choices also satisfy $\mathbf{Y}^T \mathbf{C} = \mathbf{I}$ and $\mathbf{R} \mathbf{Z} = \mathbf{I}$, but now $\mathbf{C} \mathbf{Y}^T = \mathbf{C} \mathbf{C}^I$ and $\mathbf{Z} \mathbf{R} = \mathbf{R}^I \mathbf{R}$ are *orthogonal projectors*. We compute

$$\mathbf{U} = \mathbf{Y}^T \mathbf{A} \mathbf{Z} = \mathbf{C}^I \mathbf{A} \mathbf{R}^I,$$

yielding a CUR factorization that can be viewed as a two step process: first the columns of \mathbf{A} are projected onto $\text{Ran}(\mathbf{C})$, then the result is projected onto the row space of \mathbf{R} :

$$1) \quad \mathbf{M} = \mathbf{C} \mathbf{C}^I \mathbf{A}, \quad 2) \quad \mathbf{C} \mathbf{U} \mathbf{R} = \mathbf{M} \mathbf{R}^I \mathbf{R}.$$

Both steps are optimal with respect to the 2-norm error, which is the primary source of the excellent approximation properties of this approach.

Several strategies for selecting \mathbf{p} and \mathbf{q} have been proposed.² The approach presented in the next section is simple to implement and has complexity mk and nk to select the indices \mathbf{p} and \mathbf{q} , provided the leading k right and left singular vectors of \mathbf{A} are available. Thus the overall complexity is dominated by the construction of the rank- k SVD $\mathbf{A} \approx \mathbf{V} \mathbf{S} \mathbf{W}^T$, where $\mathbf{V}^T \mathbf{V} = \mathbf{W}^T \mathbf{W} = \mathbf{I} \in \mathbb{R}^{k \times k}$ and $\mathbf{S} = \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_k)$ is the $k \times k$ matrix of dominant singular values $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_k$.

²In the theoretical computer science literature, one often takes \mathbf{C} and/or \mathbf{R} to have rank larger than k , but then builds \mathbf{U} with rank k . By selecting these extra columns and/or rows, one seeks to get within some factor $1 + \varepsilon$ of the optimal approximation; see, e.g., [5].

3 DEIM

The DEIM point selection algorithm was first presented in [6] in the context of model order reduction for nonlinear dynamical systems, and is a discrete variant of the Empirical Interpolation Method originally proposed in [4]. The DEIM procedure operates on the singular vector matrices \mathbf{V} and \mathbf{W} independently to select the row indices \mathbf{p} and column indices \mathbf{q} . We explain the process for selecting \mathbf{p} ; applying the same steps to \mathbf{W} yields \mathbf{q} . To derive the method, we elaborate upon the interpolatory projectors introduced in the last section.

Definition 3.1 *Given a full rank matrix $\mathbf{V} \in \mathbb{R}^{m \times k}$ and a set of distinct indices $\mathbf{p} \in \mathbb{N}^k$, the interpolatory projector for \mathbf{p} onto $\text{Ran}(\mathbf{V})$ is*

$$\mathcal{P} \equiv \mathbf{V}(\mathbf{P}^T \mathbf{V})^{-1} \mathbf{P}^T, \quad (3.1)$$

where $\mathbf{P} = \mathbf{I}(:, \mathbf{p}) \in \mathbb{R}^{m \times k}$, provided $\mathbf{P}^T \mathbf{V}$ is invertible.

In general \mathcal{P} is an oblique projector, and it has an important property not generally enjoyed by orthogonal projectors: for any $\mathbf{x} \in \mathbb{R}^m$,

$$(\mathcal{P}\mathbf{x})(\mathbf{p}) = \mathbf{P}^T \mathcal{P}\mathbf{x} = \mathbf{P}^T \mathbf{V}(\mathbf{P}^T \mathbf{V})^{-1} \mathbf{P}^T \mathbf{x} = \mathbf{P}^T \mathbf{x} = \mathbf{x}(\mathbf{p}),$$

so the projected vector $\mathcal{P}\mathbf{x}$ matches \mathbf{x} in the \mathbf{p} entries, justifying the name “interpolatory projector.”

The DEIM algorithm processes the columns of

$$\mathbf{V} = [\mathbf{v}_1 \quad \mathbf{v}_2 \quad \cdots \quad \mathbf{v}_k]$$

one at a time, starting from the leading singular vector \mathbf{v}_1 . Each step processes the next singular vector to produce the next index. The first index p_1 corresponds to the largest magnitude entry in \mathbf{v}_1 :

$$|\mathbf{v}_1(p_1)| = \|\mathbf{v}_1\|_\infty.$$

Now define $\mathbf{p}_1 \equiv [p_1]$, and let

$$\mathcal{P}_1 \equiv \mathbf{v}_1(\mathbf{P}_1^T \mathbf{v}_1)^{-1} \mathbf{P}_1^T$$

denote the interpolatory projector for \mathbf{p}_1 onto $\text{Ran}(\mathbf{v}_1)$. The second index p_2 corresponds to the largest entry in \mathbf{v}_2 , after the interpolatory projection in the \mathbf{v}_1 direction has been removed:

$$\begin{aligned} \mathbf{r}_2 &\equiv \mathbf{v}_2 - \mathcal{P}_1 \mathbf{v}_2 \\ |\mathbf{r}_2(p_2)| &= \|\mathbf{r}_2\|_\infty. \end{aligned}$$

Notice that $\mathbf{r}_2(p_1) = 0$, since $\mathcal{P}_1 \mathbf{v}_2$ matches \mathbf{v}_2 in the p_1 position, a consequence of interpolatory projection. This property ensures the process will never produce duplicate indices.

Now suppose we have $j - 1$ indices, with

$$\mathbf{p}_{j-1} \equiv \begin{bmatrix} p_1 \\ \vdots \\ p_{j-1} \end{bmatrix}, \quad \mathbf{P}_{j-1} \equiv \mathbf{I}(:, \mathbf{p}_{j-1}), \quad \mathbf{V}_{j-1} \equiv [\mathbf{v}_1 \quad \cdots \quad \mathbf{v}_{j-1}], \quad \mathcal{P}_{j-1} \equiv \mathbf{V}_{j-1}(\mathbf{P}_{j-1}^T \mathbf{V}_{j-1})^{-1} \mathbf{P}_{j-1}^T.$$

```

Input:  $\mathbf{V}$ , an  $m \times k$  matrix ( $m \geq k$ )
Output:  $\mathbf{p}$ , an integer vector with  $k$  distinct entries in  $\{1, \dots, m\}$ 

 $\mathbf{v} = \mathbf{V}(:, 1)$ 
 $[\sim, p_1] = \max(|\mathbf{v}|)$ 
 $\mathbf{p} = [p_1]$ 
for  $j = 2, 3, \dots, k$ 
     $\mathbf{v} = \mathbf{V}(:, j)$ 
     $\mathbf{c} = \mathbf{V}(\mathbf{p}, 1 : j - 1)^{-1} \mathbf{v}(\mathbf{p})$ 
     $\mathbf{r} = \mathbf{v} - \mathbf{V}(:, 1 : j - 1) \mathbf{c}$ 
     $[\sim, p_j] = \max(|\mathbf{r}|)$ 
     $\mathbf{p} = [\mathbf{p}; p_j]$ 
end

```

Algorithm 1: DEIM point selection algorithm.

To select p_j , remove from \mathbf{v}_j its interpolatory projection onto indices \mathbf{p}_{j-1} and take the largest remaining entry:

$$\begin{aligned} \mathbf{r}_j &\equiv \mathbf{v}_j - \mathcal{P}_{j-1} \mathbf{v}_j \\ |\mathbf{r}_j(p_j)| &= \|\mathbf{r}_j\|_\infty. \end{aligned}$$

Implementations should not explicitly construct these projectors; see the pseudocode in Algorithm 1 for details.

Those familiar with partially pivoted LU decomposition will notice, on a moment's reflection, that this index selection scheme is exactly equivalent to the index selection of partial pivoting. This arrangement is equivalent to the “left looking” variant of LU factorization [9, sect. 5.4], but with two important differences. First, there are no explicit row interchanges in DEIM, as there are in LU factorization. Second, the original basis vectors (columns of \mathbf{V}) are not replaced with the residual vectors, as happens in traditional LU decomposition. (In the context of model reduction, it is preferable to keep the nice orthogonal basis intact for use as a reduced basis.) We will exploit this connection with partially pivoted LU factorization to analyze the approximation properties of DEIM.

Since the DEIM algorithm processes the singular vectors sequentially, from most to least significant, it introduces new singular vector information in a coherent manner as it successively selects the k indices. Contrast this to index selection strategies based on leverage scores, where all singular vectors are incorporated at once via row norms of \mathbf{V} and \mathbf{W} ; to account for the fact that higher singular vectors are less significant, such approaches often instead compute leverage scores using only a few of the leading singular vectors.³

For the interpolatory projector \mathcal{P}_j to exist at the j th step, $\mathbf{P}_{j-1}^T \mathbf{V}_{j-1}$ must be nonsingular. The linear independence of the columns of \mathbf{V} assures this. In the following, \mathbf{e}_j denotes the j th column of the identity matrix.

³A potential limitation of the DEIM approach is that \mathbf{r}_j could have multiple entries that have nearly the same magnitude, but only one index is selected at the j th step; if the other large-magnitude entries in \mathbf{r}_j are not significant in subsequent \mathbf{r}_ℓ vectors, the corresponding indices will not be selected. One can imagine modifications of the selection algorithm to account for such situations, e.g., by processing multiple singular vectors at a time.

Lemma 3.1 Let $\mathbf{P}_j = [\mathbf{e}_{p_1}, \mathbf{e}_{p_2}, \dots, \mathbf{e}_{p_j}]$ and let $\mathbf{V}_j = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_j]$ for $1 \leq j \leq k$. If $\text{rank}(\mathbf{V}) = k$, then $\mathbf{P}_j^T \mathbf{V}_j$ is nonsingular for $1 \leq j \leq k$.

Proof: Suppose $\mathbf{P}_{j-1}^T \mathbf{V}_{j-1}$ is nonsingular and let $\mathbf{r}_j = \mathbf{v}_j - \mathbf{V}_{j-1}(\mathbf{P}_{j-1}^T \mathbf{V}_{j-1})^{-1} \mathbf{P}_{j-1}^T \mathbf{v}_j$. Then $\|\mathbf{r}_j\|_\infty > 0$, for otherwise $\mathbf{0} = \mathbf{v}_j - \mathbf{V}_{j-1} \mathbf{c}_{j-1}$, in violation of the assumption that $\text{rank}(\mathbf{V}) = k$. Thus

$$0 < |\mathbf{e}_{p_j}^T \mathbf{r}_j| = |\mathbf{e}_{p_j}^T \mathbf{v}_j - \mathbf{e}_{p_j}^T \mathbf{V}_{j-1}(\mathbf{P}_{j-1}^T \mathbf{V}_{j-1})^{-1} \mathbf{P}_{j-1}^T \mathbf{v}_j|, \quad (3.2)$$

where p_j is the j th DEIM interpolation point. Now factor

$$\mathbf{P}_j^T \mathbf{V}_j = \begin{bmatrix} \mathbf{P}_{j-1}^T \mathbf{V}_{j-1} & \mathbf{P}_{j-1}^T \mathbf{v}_j \\ \mathbf{e}_{p_j}^T \mathbf{V}_{j-1} & \mathbf{e}_{p_j}^T \mathbf{v}_j \end{bmatrix} = \begin{bmatrix} \mathbf{I}_{j-1} & \mathbf{0} \\ \mathbf{e}_{p_j}^T \mathbf{V}_{j-1}(\mathbf{P}_{j-1}^T \mathbf{V}_{j-1})^{-1} & 1 \end{bmatrix} \begin{bmatrix} \mathbf{P}_{j-1}^T \mathbf{V}_{j-1} & \mathbf{P}_{j-1}^T \mathbf{v}_j \\ \mathbf{0} & \nu_j \end{bmatrix}, \quad (3.3)$$

where

$$\nu_j = \mathbf{e}_{p_j}^T \mathbf{v}_j - \mathbf{e}_{p_j}^T \mathbf{V}_{j-1}(\mathbf{P}_{j-1}^T \mathbf{V}_{j-1})^{-1} \mathbf{P}_{j-1}^T \mathbf{v}_j.$$

The inequality (3.2) implies $\nu_j \neq 0$ and hence equation (3.3) implies $\mathbf{P}_j^T \mathbf{V}_j$ is nonsingular. Since $\mathbf{e}_{p_1}^T \mathbf{v}_1 \neq 0$, this argument provides an inductive proof that $\mathbf{P}_j^T \mathbf{V}_j$ is nonsingular for $1 \leq j \leq k$. \blacksquare

4 CUR Approximation Properties

While the theory presented in this section was designed to bound $\|\mathbf{A} - \mathbf{CUR}\|$ for the DEIM-CUR method, the analysis applies to *any* CUR factorization with full rank $\mathbf{C} \in \mathbb{R}^{m \times k}$ and $\mathbf{R} \in \mathbb{R}^{k \times n}$, and $\mathbf{U} = \mathbf{C}^I \mathbf{A} \mathbf{R}^I$, regardless of the procedure used for selecting the columns and rows.⁴

Consider a CUR factorization that uses row indices $\mathbf{p} \in \mathbb{N}^k$ and column indices $\mathbf{q} \in \mathbb{N}^k$, and set

$$\mathbf{P} = \mathbf{I}(:, \mathbf{p}) = [\mathbf{e}_{p_1}, \dots, \mathbf{e}_{p_k}] \in \mathbb{R}^{m \times k}, \quad \mathbf{Q} = \mathbf{I}(:, \mathbf{q}) = [\mathbf{e}_{q_1}, \dots, \mathbf{e}_{q_k}] \in \mathbb{R}^{n \times k}.$$

The first step in this analysis bounds the mismatch between \mathbf{A} and its interpolatory projection $\mathcal{P}\mathbf{A}$.

Lemma 4.1 Assume $\mathbf{P}^T \mathbf{V}$ is invertible and let $\mathcal{P} = \mathbf{V}(\mathbf{P}^T \mathbf{V})^{-1} \mathbf{P}^T$ be the interpolatory projector (3.1). If $\mathbf{V}^T \mathbf{V} = \mathbf{I}$, then any $\mathbf{A} \in \mathbb{R}^{m \times n}$ satisfies

$$\|\mathbf{A} - \mathcal{P}\mathbf{A}\| \leq \|(\mathbf{P}^T \mathbf{V})^{-1}\| \|(\mathbf{I} - \mathbf{V}\mathbf{V}^T)\mathbf{A}\|.$$

Additionally, if \mathbf{V} consists of the leading k left singular vectors of \mathbf{A} , then

$$\|\mathbf{A} - \mathcal{P}\mathbf{A}\| = \|(\mathbf{I} - \mathcal{P})\mathbf{A}\| \leq \|(\mathbf{P}^T \mathbf{V})^{-1}\| \sigma_{k+1}.$$

⁴We are grateful to Ilse Ipsen for noting the applicability of this analysis to all such **CUR** factorizations, and for also pointing out that, given knowledge of *all* the singular values and vectors of \mathbf{A} , our Lemma 4.2 can be sharpened via application of [16, Thm. 9.1]. Indeed, Ipsen observes that the interpolatory projector proof of Lemma 4.2 can be adapted to simplify the multipage proof of [16, Thm. 9.1].

Proof: First note that $\mathcal{P}\mathbf{V} = \mathbf{V}(\mathbf{P}^T\mathbf{V})^{-1}\mathbf{P}^T\mathbf{V} = \mathbf{V}$, so that $(\mathbf{I} - \mathcal{P})\mathbf{V} = \mathbf{0}$. Therefore

$$\|\mathbf{A} - \mathcal{P}\mathbf{A}\| = \|(\mathbf{I} - \mathcal{P})\mathbf{A}\| = \|(\mathbf{I} - \mathcal{P})(\mathbf{I} - \mathbf{V}\mathbf{V}^T)\mathbf{A}\| \leq \|(\mathbf{I} - \mathcal{P})\| \|(\mathbf{I} - \mathbf{V}\mathbf{V}^T)\mathbf{A}\|.$$

It is well known that

$$\|\mathbf{I} - \mathcal{P}\| = \|\mathcal{P}\| = \|(\mathbf{P}^T\mathbf{V})^{-1}\|$$

so long as $\mathcal{P} \neq \mathbf{0}$ or \mathbf{I} ; see, e.g., [24]. This establishes the first result. The second follows from the fact that

$$\|(\mathbf{I} - \mathbf{V}\mathbf{V}^T)\mathbf{A}\| = \|\mathbf{A} - \mathbf{V}\mathbf{S}\mathbf{W}^T\| = \sigma_{k+1}$$

when \mathbf{V} consists of the leading k left singular vectors of \mathbf{A} . ■

Now let $\mathbf{V}\mathbf{S}\mathbf{W}^T \approx \mathbf{A}$ be a rank- k SVD of \mathbf{A} . (The singular vectors play a crucial role in this analysis, even if \mathbf{p} and \mathbf{q} were selected using some scheme that did not reference them.) In addition to the interpolatory projector $\mathcal{P} = \mathbf{V}(\mathbf{P}^T\mathbf{V})^{-1}\mathbf{P}^T$ that operates on the left of \mathbf{A} , we shall also use $\mathcal{Q} = \mathbf{Q}(\mathbf{W}^T\mathbf{Q})^{-1}\mathbf{W}^T$, which operates on the right of \mathbf{A} . Assuming that $\mathbf{P}^T\mathbf{V}$ and $\mathbf{W}^T\mathbf{Q}$ are invertible, define the error constants

$$\eta_p \equiv \|(\mathbf{P}^T\mathbf{V})^{-1}\|, \quad \eta_q \equiv \|(\mathbf{W}^T\mathbf{Q})^{-1}\|.$$

Lemma 4.1 implies

$$\|\mathbf{A}(\mathbf{I} - \mathcal{Q})\| \leq \eta_q \sigma_{k+1} \quad \text{and} \quad \|(\mathbf{I} - \mathcal{P})\mathbf{A}\| \leq \eta_p \sigma_{k+1}. \quad (4.1)$$

The next lemma shows that these bounds on the error of the interpolatory projection of \mathbf{A} onto the select columns and rows also apply to the *orthogonal* projections of \mathbf{A} onto the same column and row spaces.

Lemma 4.2 *Suppose the row and column indices \mathbf{p} and \mathbf{q} give full rank matrices $\mathbf{C} = \mathbf{A}(:, \mathbf{q}) = \mathbf{A}\mathbf{Q} \in \mathbb{R}^{m \times k}$ and $\mathbf{R} = \mathbf{A}(\mathbf{p}, :) = \mathbf{P}\mathbf{A} \in \mathbb{R}^{k \times n}$, with finite error constants η_p and η_q , and suppose that $k < \min\{m, n\}$. Then*

$$\|(\mathbf{I} - \mathbf{C}\mathbf{C}^I)\mathbf{A}\| \leq \eta_q \sigma_{k+1} \quad \text{and} \quad \|\mathbf{A}(\mathbf{I} - \mathbf{R}^I\mathbf{R})\| \leq \eta_p \sigma_{k+1}.$$

Proof: Using the formula $\mathbf{C} = \mathbf{A}\mathbf{Q}$, we have $\mathbf{C}^I = (\mathbf{C}^T\mathbf{C})^{-1}\mathbf{C}^T = (\mathbf{Q}^T\mathbf{A}^T\mathbf{A}\mathbf{Q})^{-1}(\mathbf{A}\mathbf{Q})^T$, so the orthogonal projection of \mathbf{A} onto $\text{Ran}(\mathbf{C})$ is

$$\mathbf{C}\mathbf{C}^I\mathbf{A} = (\mathbf{A}\mathbf{Q}(\mathbf{Q}^T\mathbf{A}^T\mathbf{A}\mathbf{Q})^{-1}\mathbf{Q}^T\mathbf{A}^T)\mathbf{A} = \mathbf{A}(\mathbf{Q}(\mathbf{Q}^T\mathbf{A}^T\mathbf{A}\mathbf{Q})^{-1}\mathbf{Q}^T\mathbf{A}^T\mathbf{A}).$$

Hence the error in the orthogonal projection of \mathbf{A} is

$$(\mathbf{I} - \mathbf{C}\mathbf{C}^I)\mathbf{A} = \mathbf{A}(\mathbf{I} - \mathbf{\Phi}), \quad \text{where} \quad \mathbf{\Phi} = \mathbf{Q}(\mathbf{Q}^T\mathbf{A}^T\mathbf{A}\mathbf{Q})^{-1}\mathbf{Q}^T\mathbf{A}^T\mathbf{A}. \quad (4.2)$$

Note that $\mathbf{\Phi}$ is an oblique projector onto $\text{Ran}(\mathbf{Q})$, so $\mathbf{\Phi}\mathbf{Q} = \mathbf{Q}$. Therefore, $\mathbf{\Phi}\mathcal{Q} = \mathcal{Q}$, since

$$\mathbf{\Phi}\mathcal{Q} = \mathbf{\Phi}\mathbf{Q}(\mathbf{W}^T\mathbf{Q})^{-1}\mathbf{W}^T = \mathbf{Q}(\mathbf{W}^T\mathbf{Q})^{-1}\mathbf{W}^T = \mathcal{Q}.$$

This implies that

$$\mathbf{A}(\mathbf{I} - \mathbf{\Phi}) = \mathbf{A}(\mathbf{I} - \mathbf{\Phi})(\mathbf{I} - \mathcal{Q}) = (\mathbf{I} - \mathbf{C}\mathbf{C}^I)\mathbf{A}(\mathbf{I} - \mathcal{Q}),$$

and so from (4.2) we have

$$\begin{aligned}
\|(\mathbf{I} - \mathbf{C}\mathbf{C}^I)\mathbf{A}\| &= \|\mathbf{A}(\mathbf{I} - \mathbf{\Phi})\| \\
&= \|(\mathbf{I} - \mathbf{C}\mathbf{C}^I)\mathbf{A}(\mathbf{I} - \mathbf{Q})\| \\
&\leq \|\mathbf{I} - \mathbf{C}\mathbf{C}^I\| \|\mathbf{A}(\mathbf{I} - \mathbf{Q})\| \\
&\leq \eta_q \sigma_{k+1}.
\end{aligned}$$

The last line follows from the bound (4.1) and the fact that $\|\mathbf{I} - \mathbf{C}\mathbf{C}^I\| = 1$, since $\mathbf{C}\mathbf{C}^I$ is an orthogonal projector and $k < \min\{m, n\}$.

A similar argument shows that

$$\mathbf{A}(\mathbf{I} - \mathbf{R}^I\mathbf{R}) = (\mathbf{I} - \mathbf{\Psi})\mathbf{A}$$

where $\mathbf{\Psi} = \mathbf{A}\mathbf{A}^T\mathbf{P}(\mathbf{P}^T\mathbf{A}\mathbf{A}^T\mathbf{P})^{-1}\mathbf{P}^T$, and also that

$$(\mathbf{I} - \mathbf{\Psi})\mathbf{A} = (\mathbf{I} - \mathcal{P})(\mathbf{I} - \mathbf{\Psi})\mathbf{A} = (\mathbf{I} - \mathcal{P})\mathbf{A}(\mathbf{I} - \mathbf{R}^I\mathbf{R}),$$

from which follows the error bound

$$\|\mathbf{A}(\mathbf{I} - \mathbf{R}^I\mathbf{R})\| \leq \|(\mathbf{I} - \mathcal{P})\mathbf{A}\| \|\mathbf{I} - \mathbf{R}^I\mathbf{R}\| \leq \eta_p \sigma_{k+1}. \quad \blacksquare$$

The main result on approximation of \mathbf{A} by \mathbf{CUR} readily follows from combining this last lemma with a basic CUR analysis technique used by Mahoney and Drineas [22, eq. (6)].

Theorem 4.1 *Given $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $1 \leq k < \min\{m, n\}$, let $\mathbf{C} = \mathbf{A}(:, \mathbf{q}) \in \mathbb{R}^{m \times k}$ and $\mathbf{R} = \mathbf{A}(\mathbf{p}, :) \in \mathbb{R}^{k \times n}$ with finite error constants η_p and η_q , and set $\mathbf{U} = \mathbf{C}^I\mathbf{A}\mathbf{R}^I$. Then*

$$\|\mathbf{A} - \mathbf{CUR}\| \leq (\eta_p + \eta_q) \sigma_{k+1}.$$

Proof: From the definitions,

$$\mathbf{A} - \mathbf{CUR} = \mathbf{A} - \mathbf{C}\mathbf{C}^I\mathbf{A}\mathbf{R}^I\mathbf{R} = (\mathbf{I} - \mathbf{C}\mathbf{C}^I)\mathbf{A} + \mathbf{C}\mathbf{C}^I\mathbf{A}(\mathbf{I} - \mathbf{R}^I\mathbf{R}).$$

Applying Lemma 4.2,

$$\begin{aligned}
\|\mathbf{A} - \mathbf{CUR}\| &\leq \|(\mathbf{I} - \mathbf{C}\mathbf{C}^I)\mathbf{A}\| + \|\mathbf{C}\mathbf{C}^I\| \|\mathbf{A}(\mathbf{I} - \mathbf{R}^I\mathbf{R})\| \\
&\leq \eta_q \sigma_{k+1} + \eta_p \sigma_{k+1} \\
&= (\eta_p + \eta_q) \sigma_{k+1},
\end{aligned}$$

since $\|\mathbf{C}\mathbf{C}^I\| = 1$. \blacksquare

Theorem 4.1 shows that \mathbf{CUR} is within a factor of $\eta_p + \eta_q$ of the optimal rank- k approximation, hence these error constants suggest a way to assess a wide variety of column/row selection schemes. The quality of the approximation is controlled by the conditioning of the selected k rows of the dominant k (exact) singular vectors. If those singular vectors are available as part of the column/row selection process, then Theorem 4.1 provides an

a *posteriori* bound requiring only the fast ($\mathcal{O}(k^3)$) computation of η_p and η_q , and thus could suggest methods for adjusting either k or the point selection process to reduce the error constants. In this context, notice that if $\mathbf{V}\mathbf{S}\mathbf{W}^T$ is only an *approximation* to the optimal rank- k SVD with \mathbf{V} and \mathbf{W} having orthonormal columns (as computed, for example, using the incremental QR algorithm described in the next section), the preceding analysis gives

$$\begin{aligned}\|\mathbf{A} - \mathbf{CUR}\| &\leq \|(\mathbf{I} - \mathbf{C}\mathbf{C}^T)\mathbf{A}\| + \|\mathbf{A}(\mathbf{I} - \mathbf{R}^T\mathbf{R})\| \\ &= \|\mathbf{A}(\mathbf{I} - \mathbf{Q})\| + \|(\mathbf{I} - \mathbf{P})\mathbf{A}\| \\ &\leq \|(\mathbf{W}^T\mathbf{Q})^{-1}\|\|\mathbf{A}(\mathbf{I} - \mathbf{W}\mathbf{W}^T)\| + \|(\mathbf{P}^T\mathbf{V})^{-1}\|\|(\mathbf{I} - \mathbf{V}\mathbf{V}^T)\mathbf{A}\|,\end{aligned}\tag{4.3}$$

showing how σ_{k+1} in Theorem 4.1 is replaced by the error in the approximate SVD through $\|\mathbf{A}(\mathbf{I} - \mathbf{W}\mathbf{W}^T)\|$ and $\|(\mathbf{I} - \mathbf{V}\mathbf{V}^T)\mathbf{A}\|$. In this case $\|(\mathbf{W}^T\mathbf{Q})^{-1}\|$ and $\|(\mathbf{P}^T\mathbf{V})^{-1}\|$ are computed using the approximate singular vectors in \mathbf{V} and \mathbf{W} , rather than the exact singular vectors in the theorem. Alternatively, if one has probabilistic bounds for η_p and η_q , then Theorem 4.1 immediately gives a probabilistic bound for $\|\mathbf{A} - \mathbf{CUR}\|$.

Numerical examples in Section 6 compare how the error constants evolve as k increases for the DEIM-CUR factorization and several other factorizations based on leverage scores.

4.1 Interpretation of the bound for DEIM-CUR

For DEIM-CUR, we can ensure the hypotheses of Theorem 4.1 are satisfied and bound the error constants. Suppose the DEIM points are selected using the exact rank- k SVD $\mathbf{A} \approx \mathbf{V}\mathbf{S}\mathbf{W}^T$. Lemma 3.1 ensures that the matrices $\mathbf{P}^T\mathbf{V}$ and $\mathbf{W}^T\mathbf{Q}$ are invertible, so η_p and η_q are finite. The DEIM strategy also gives full rank \mathbf{C} and \mathbf{R} matrices, presuming $k \leq \text{rank}(\mathbf{A})$. To see this, note that for any unit vector $\mathbf{y} \in \mathbb{R}^k$,

$$\mathbf{C}\mathbf{y} = \mathbf{A}\mathbf{Q}\mathbf{y} = \mathbf{V}\mathbf{S}\mathbf{W}^T\mathbf{Q}\mathbf{y} + \mathbf{E}\mathbf{Q}\mathbf{y},$$

where $\mathbf{E} = \mathbf{A} - \mathbf{V}\mathbf{S}\mathbf{W}^T$. Since $\mathbf{V}^T\mathbf{E} = \mathbf{0}$,

$$\|\mathbf{C}\mathbf{y}\|^2 = \|\mathbf{A}\mathbf{Q}\mathbf{y}\|^2 = \|\mathbf{V}\mathbf{S}\mathbf{W}^T\mathbf{Q}\mathbf{y}\|^2 + \|\mathbf{E}\mathbf{Q}\mathbf{y}\|^2.$$

Since $\|\mathbf{W}^T\mathbf{Q}\mathbf{y}\| \geq \|\mathbf{y}\|/\|(\mathbf{W}^T\mathbf{Q})^{-1}\| = 1/\eta_q$,

$$\|\mathbf{C}\mathbf{y}\| \geq \|\mathbf{V}\mathbf{S}\mathbf{W}^T\mathbf{Q}\mathbf{y}\| \geq \sigma_k/\eta_q > 0.$$

Thus \mathbf{C} must be full rank. A similar argument shows \mathbf{R} to be full rank as well.

The examples in Section 6 illustrate that η_p and η_q are often quite modest for the DEIM-CUR approach, e.g., $\mathcal{O}(100)$. However, worst-case bounds permit significant growth in k that is generally not observed in practice. We begin by stating a bound on this growth developed by Chaturantabut and Sorensen [6, Lemma 3.2].

Lemma 4.3 *For the DEIM selection scheme derived above,*

$$\eta_p \leq \frac{(1 + \sqrt{2m})^{k-1}}{\|\mathbf{v}_1\|_\infty}, \quad \eta_q \leq \frac{(1 + \sqrt{2n})^{k-1}}{\|\mathbf{w}_1\|_\infty},$$

where \mathbf{v}_1 and \mathbf{w}_1 denote the first columns of \mathbf{V} and \mathbf{W} .

Motivated by recent work by Drmač and Gugercin [11] on a modified DEIM-like algorithm for model reduction, we can improve this bound considerably.

Lemma 4.4 *For the DEIM selection scheme derived above,*

$$\eta_p < \sqrt{\frac{mk}{3}} 2^k, \quad \eta_q < \sqrt{\frac{nk}{3}} 2^k.$$

Proof: We shall prove the result for η_p ; the result for η_q follows similarly. As usual, let $\mathbf{V} \in \mathbb{R}^{m \times k}$ have orthonormal columns, and let $\mathbf{p} = \text{DEIM}(\mathbf{V})$ denote the row index vector derived from the DEIM selection scheme described above. Let $\mathbf{P} = \mathbf{I}(:, \mathbf{p})$ so that $\mathbf{P}^T \mathbf{V} = \mathbf{V}(\mathbf{p}, :)$.

Without loss of generality, assume the DEIM index selection gives $\mathbf{p} = [1, 2, \dots, k]^T$. (Otherwise, introduce a permutation matrix to the argument that follows.) As described in section 3, the DEIM index selection is precisely the index selection of LU decomposition with partial pivoting, so one can write

$$\mathbf{V} = \mathbf{L}\mathbf{T},$$

where the nonsingular matrix $\mathbf{T} \in \mathbb{R}^{k \times k}$ is upper triangular and $\mathbf{L} \in \mathbb{R}^{m \times k}$ is unit lower triangular with $|\mathbf{L}(i, j)| \leq 1$, $\mathbf{L}(j, j) = 1$, $1 \leq j \leq k$ and $\mathbf{L}(i, j) = 0$, $j > i$.

Let $\mathbf{L}_1 \equiv \mathbf{L}(1 : k, 1 : k)$. Then $\mathbf{V}(\mathbf{p}, :) = \mathbf{L}_1 \mathbf{T}$ and thus

$$\eta_p \equiv \|(\mathbf{P}^T \mathbf{V})^{-1}\| = \|(\mathbf{L}_1 \mathbf{T})^{-1}\| \leq \|\mathbf{T}^{-1}\| \|\mathbf{L}_1^{-1}\|.$$

(The linear independence of the columns of \mathbf{V} ensure that \mathbf{L}_1 and \mathbf{T} are invertible.) Upper bounds for $\|\mathbf{T}^{-1}\|$ and $\|\mathbf{L}_1^{-1}\|$ will give an upper bound for η_p .

To bound $\|\mathbf{T}^{-1}\|$, let $\mathbf{y} \in \mathbb{R}^k$ be a unit vector such that $\|\mathbf{T}^{-1} \mathbf{y}\| = \|\mathbf{T}^{-1}\|$. Then

$$\|\mathbf{T}^{-1}\| = \|\mathbf{T}^{-1} \mathbf{y}\| = \|\mathbf{V} \mathbf{T}^{-1} \mathbf{y}\| = \|\mathbf{L} \mathbf{y}\|.$$

Now

$$\|\mathbf{L} \mathbf{y}\| \leq \sqrt{m} \|\mathbf{L} \mathbf{y}\|_\infty = \sqrt{m} |\mathbf{e}_j^T \mathbf{L} \mathbf{y}|,$$

for some index $j \in \{1, \dots, m\}$. By the Cauchy–Schwarz inequality and the bound $|\mathbf{L}(i, j)| \leq 1$,

$$|\mathbf{e}_j^T \mathbf{L} \mathbf{y}| \leq \|\mathbf{L}^T \mathbf{e}_j\| \|\mathbf{y}\| \leq \sqrt{k} \cdot 1,$$

and so it follows that $\|\mathbf{T}^{-1}\| \leq \sqrt{mk}$.

The inverse of \mathbf{L}_1 can be bounded using forward substitution. Let $\mathbf{L}_1 \mathbf{z} = \mathbf{y}$, where $\|\mathbf{y}\| = 1$ and $\|\mathbf{z}\| = \|\mathbf{L}_1^{-1}\|$. Forward substitution provides

$$\begin{aligned} \zeta_1 &= \gamma_1 \\ \zeta_i &= \gamma_i - \sum_{j=1}^{i-1} \lambda_{ij} \zeta_j, \quad i = 2, \dots, k, \end{aligned}$$

where $\zeta_i = \mathbf{z}(i)$, $\gamma_i = \mathbf{y}(i)$ and $\lambda_{ij} = \mathbf{L}(i, j)$. We now use induction to prove

$$|\zeta_i| \leq 2^{i-1}, \quad 1 \leq i \leq k.$$

First note that $\zeta_1 = \gamma_1$, so $|\zeta_1| \leq |\gamma_1| \leq 1 = 2^0$ to establish the base case. Assume for some $i \geq 1$ that

$$|\zeta_j| \leq 2^{j-1}, \quad 1 \leq j \leq i.$$

Then

$$\begin{aligned} |\zeta_{i+1}| &= \left| \gamma_{i+1} - \sum_{j=1}^i \lambda_{ij} \zeta_j \right| \leq |\gamma_{i+1}| + \sum_{j=1}^i |\lambda_{ij}| |\zeta_j| \\ &\leq 1 + \sum_{j=1}^i 1 \cdot 2^{j-1} = 1 + \sum_{j=0}^{i-1} 2^j = 1 + (2^i - 1) = 2^i \end{aligned}$$

to complete the induction. Now since $\|\mathbf{z}\| = \|\mathbf{L}_1^{-1}\|$,

$$\|\mathbf{L}_1^{-1}\|^2 = \mathbf{z}^T \mathbf{z} = \sum_{i=1}^k |\zeta_i|^2 \leq \sum_{i=0}^{k-1} 4^i = (4^k - 1)/3.$$

Thus $\|\mathbf{L}_1^{-1}\| < 2^k/\sqrt{3}$, which, together with the bound on the inverse of \mathbf{T} , provides the final result for $m > k$:

$$\eta_p \equiv \|(\mathbf{P}^T \mathbf{V})^{-1}\| < \sqrt{\frac{mk}{3}} 2^k.$$

If $m = k$, then $\eta_p = 1$, and the result holds trivially. ■

Note that this proof only relies on the orthonormality of the columns of \mathbf{V} and \mathbf{W} , and hence it applies when the DEIM selection scheme is applied to approximate singular vectors, as in CUR error bound in (4.3).

Lemma 4.4 was inspired by the proof technique developed by Drmač and Gugercin [11] to bound $\|(\mathbf{P}^T \mathbf{V})^{-1}\|$, when \mathbf{P} is selected by applying a pivoted rank-revealing QR factorization scheme to \mathbf{V} . Note that this new bound is on the same order of magnitude as the Drmač–Gugercin scheme. In practice, their scheme seems to give slightly smaller growth that is more consistent over a wide range of examples. Neither scheme experienced exponential growth over very extensive testing. For the DEIM approach, this absence of exponential growth is closely related to decades of experience with Gaussian elimination with partial pivoting. Element growth in \mathbf{T} is bounded by a factor of 2^{k-1} (for a $k \times k$ matrix), and there is an example that achieves this growth. Nevertheless, this algorithm is almost exclusively used to solve linear systems because such growth is never experienced.⁵ Indeed, a similar near worst case example can be constructed for DEIM, although this growth has not been observed in practice.

⁵See, for example, the extensive numerical tests involving random matrices described in [26, lecture 22] and [27]. Interestingly, in the experiments of Trefethen and Schreiber [27], random matrices with orthonormal columns tend to have slightly larger growth factors than Gaussian matrices, though both cases are very far indeed from the exponential upper bound.

A Growth Example: We now construct an orthonormal matrix \mathbf{V} with the property

$$\frac{1}{\sqrt{8}} 2^k < \eta_p \equiv \|(\mathbf{P}^T \mathbf{V})^{-1}\| < \sqrt{\frac{mk}{3}} 2^k \quad (4.4)$$

where $\mathbf{P}^T \mathbf{V} = \mathbf{V}(\mathbf{p}, :)$ with $\mathbf{p} = \text{DEIM}(\mathbf{V})$. To construct \mathbf{V} , begin by defining

$$\mathbf{L} := \begin{bmatrix} 1 & & & \\ -1 & 1 & & \\ \vdots & \ddots & \ddots & \\ -1 & \cdots & -1 & 1 \\ -1 & \cdots & -1 & -1 \\ \vdots & & & \vdots \\ -1 & \cdots & -1 & -1 \end{bmatrix} \in \mathbb{R}^{m \times k}.$$

Now construct $\mathbf{V}\mathbf{T}_1 \equiv \mathbf{L}$ as an economy-sized QR factorization of \mathbf{L} (with no column pivoting). Since the columns of \mathbf{L} are linearly independent by construction, $\mathbf{T}_1 \in \mathbb{R}^{k \times k}$ is invertible; define $\mathbf{T} \equiv \mathbf{T}_1^{-1}$, so that $\mathbf{V} = \mathbf{L}\mathbf{T}$. (Note that \mathbf{T} plays the same role it does in the proof of Lemma 4.4.) If the DEIM procedure is applied to \mathbf{V} , then by construction $\mathbf{p} = [1, 2, \dots, k]$ (in exact arithmetic): during the DEIM procedure, the relations

$$\ell_j \tau_{jj} = \mathbf{v}_j - \mathbf{V}_{j-1}(\mathbf{P}_{j-1}^T \mathbf{V}_{j-1})^{-1} \mathbf{P}_{j-1}^T \mathbf{v}_j, \quad j > 1$$

hold, with $\ell_j = \mathbf{L}(:, j)$, $\tau_{jj} = \mathbf{T}(j, j)$, $\mathbf{v}_j = \mathbf{V}(:, j)$, $\mathbf{P}_{j-1} = \mathbf{I}(:, 1:j-1)$ and $\mathbf{V}_{j-1} = \mathbf{V}(:, 1:j-1)$. Thus $\mathbf{p}(j) = j$, $j > 1$ and it is easily seen that $\mathbf{p}(1) = 1$.

Note that $\mathbf{V}\mathbf{T}^{-1} = \mathbf{L}$ implies $\mathbf{T}^{-T}\mathbf{T}^{-1} = \mathbf{L}^T\mathbf{L}$ hence $\|\mathbf{T}\| = 1/\sigma_k$, where σ_k is the smallest singular value of \mathbf{L} . Let \mathbf{y} be the corresponding right singular vector, so that

$$\sigma_k^2 = \mathbf{y}^T \mathbf{L}^T \mathbf{L} \mathbf{y}.$$

We claim that $\sigma_k \geq \sqrt{2}$. To see this, write \mathbf{L} in the form

$$\mathbf{L} = \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0} \end{bmatrix} - \begin{bmatrix} \mathbf{L}_0 \\ \mathbf{E} \end{bmatrix},$$

where

$$\mathbf{I}_k = \begin{bmatrix} 1 & & & \\ & 1 & & \\ & & \ddots & \\ & & & 1 \end{bmatrix}, \quad \mathbf{L}_0 = \begin{bmatrix} 0 & & & \\ 1 & \ddots & & \\ \vdots & \ddots & \ddots & \\ 1 & \cdots & 1 & 0 \end{bmatrix}, \quad \mathbf{E} = \begin{bmatrix} 1 & 1 & \cdots & 1 \\ 1 & 1 & \cdots & 1 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \cdots & 1 \end{bmatrix} = \mathbf{f} \mathbf{e}^T,$$

for $\mathbf{f} = [1, \dots, 1]^T \in \mathbb{R}^{m-k}$.

$$\begin{aligned} \mathbf{L}^T \mathbf{L} &= \mathbf{I}_k - \mathbf{L}_0 - \mathbf{L}_0^T + \mathbf{L}_0^T \mathbf{L}_0 + \mathbf{E}^T \mathbf{E} \\ &= \mathbf{I}_k - (\mathbf{e} \mathbf{e}^T - \mathbf{I}_k) + \mathbf{L}_0^T \mathbf{L}_0 + (m-k) \mathbf{e} \mathbf{e}^T \\ &= 2\mathbf{I}_k + \mathbf{L}_0^T \mathbf{L}_0 + (m-k-1) \mathbf{e} \mathbf{e}^T \\ &= 2\mathbf{I}_k + \mathbf{M}, \end{aligned}$$

where $\mathbf{M} := \mathbf{L}_0^T \mathbf{L}_0 + (m - k - 1)\mathbf{e}\mathbf{e}^T$ is symmetric and positive semidefinite whenever $m > k$. Thus

$$\sigma_k^2 = \mathbf{y}^T \mathbf{L}^T \mathbf{L} \mathbf{y} = \mathbf{y}^T (2\mathbf{I}_k + \mathbf{M}) \mathbf{y} \geq 2$$

and hence it follows that

$$\|\mathbf{T}\| \leq 1/\sqrt{2}.$$

This implies

$$\|(\mathbf{P}^T \mathbf{V})^{-1}\| = \|\mathbf{T}^{-1} \mathbf{L}_1^{-1}\| \geq \frac{\|\mathbf{L}_1^{-1}\|}{\|\mathbf{T}\|} \geq \sqrt{2} \|\mathbf{L}_1^{-1}\|.$$

To complete the lower bound, we must analyze $\|\mathbf{L}_1^{-1}\|$. Forward substitution gives $\mathbf{L}_1^{-1} \mathbf{e}_1 = [1, 1, 2, 4, \dots, 2^{k-2}]^T$ and thus

$$\|\mathbf{L}_1^{-1}\| > \|\mathbf{L}_1^{-1} \mathbf{e}_1\| = \sqrt{1 + (4^{k-1} - 1)/3} > 2^{k-2}.$$

We arrive at the lower bound

$$\eta_p \equiv \|(\mathbf{P}^T \mathbf{V})^{-1}\| \geq \sqrt{2} \|\mathbf{L}_1^{-1}\| > \sqrt{2} \cdot 2^{k-2},$$

and thus for this choice of \mathbf{V} , the DEIM error constant satisfies

$$\frac{1}{\sqrt{8}} 2^k < \eta_p < \sqrt{\frac{mk}{3}} 2^k.$$

This example is interesting because it relies on the behavior of the classic example for growth in LU decomposition [26, lecture 22]. However, in this case the pathological growth is caused by \mathbf{L} and not by \mathbf{T} .

5 Incremental QR Factorization

The DEIM point selection process presumes access to the first k left and right singular vectors of $\mathbf{A} \in \mathbb{R}^{m \times n}$. If either m or n is of modest size (say ≤ 1000) and \mathbf{A} can be stored as a dense matrix, library software for computing the “economy sized” SVD, e.g., `[V, S, W] = svd(A, 'econ')` in MATLAB, usually performs very well. For larger scale problems, the leading k singular vectors can be computed using iterative methods, such as the Krylov subspace-based ARPACK software [21] (used by MATLAB’s `svds` command), PROPACK [20], IRLBA [1], or the Jacobi–Davidson algorithm [17]. Randomized SVD algorithms provide an appealing alternative with probabilistic error bounds [16]. Here we describe another approach that satisfies a deterministic error bound (Lemma 5.1) and *only requires one pass through the matrix \mathbf{A}* , a key property for massive data sets that cannot easily be stored in memory.

This approach is based on an *incremental* low rank $\mathbf{A} \approx \mathbf{QR}$ approximation, where $\mathbf{Q} \in \mathbb{R}^{n \times k}$ has orthonormal columns and $\mathbf{R} \in \mathbb{R}^{k \times m}$ is upper triangular. (In this section only, \mathbf{Q} and \mathbf{R} denote different quantities from elsewhere in the paper.) Take the dense (economy sized) SVD $\mathbf{R} = \hat{\mathbf{V}} \mathbf{S} \mathbf{W}^T$, and put $\mathbf{V} = \mathbf{Q} \hat{\mathbf{V}}$ to get

$$\mathbf{A} \approx \mathbf{QR} = \mathbf{V} \mathbf{S} \mathbf{W}^T. \tag{5.1}$$

Input: \mathbf{A} , an $m \times n$ matrix
 tol , a positive scalar controlling the accuracy of the factorization

Output: \mathbf{Q} , an $m \times k$ matrix with orthonormal columns
 \mathbf{R} , a $k \times n$ rectangular matrix
with $\mathbf{A} \approx \mathbf{QR}$

Choose $k \ll \min(m, n)$
Compute the QR factorization $\mathbf{A}(:, 1:k) = \mathbf{QR}$, with $\mathbf{Q} \in \mathbb{R}^{n \times k}$ and $\mathbf{R} \in \mathbb{R}^{k \times m}$
 $rownorms(i) = \|\mathbf{R}(i, :)\|^2$ for $i = 1, \dots, k$
 $j = k + 1$
while $j \leq n$
 $\mathbf{a} = \mathbf{A}(:, j)$; $\mathbf{r} = \mathbf{Q}^T \mathbf{a}$; $\mathbf{f} = \mathbf{a} - \mathbf{Q}\mathbf{r}$; $\rho = \|\mathbf{f}\|$; $\mathbf{q} = \mathbf{f}/\rho$
 $\mathbf{Q} = [\mathbf{Q}, \mathbf{q}]$; $\mathbf{R} = \begin{bmatrix} \mathbf{R} & \mathbf{r} \\ \mathbf{0} & \rho \end{bmatrix}$
 $rownorms(i) = rownorms(i) + \mathbf{r}(i)^2$ for $i = 1, \dots, k$
 $rownorms(k+1) = \rho^2$;
 $FnormR = \text{sum}(rownorms)$;
 $[\sigma, i_{\min}] = \min(rownorms(1:k+1))$;
 if $\sigma > (tol^2) * (FnormR - rownorms(i_{\min}))$
 % no deflation
 $k = k + 1$;
 else *% deflation required*
 if $i_{\min} < k + 1$
 $\mathbf{R}(i_{\min}, :) = \mathbf{R}(k+1, :)$; $\mathbf{Q}(:, i_{\min}) = \mathbf{Q}(:, k+1)$
 $rownorms(i_{\min}) = rownorms(k+1)$
 end
 % delete the minimum norm row of R
 $\mathbf{Q} = \mathbf{Q}(:, 1:k)$; $\mathbf{R} = \mathbf{R}(1:k, :)$
 end
 $j = j + 1$
end

Algorithm 2: Incremental QR low rank approximate factorization

Incremental algorithms for building the QR factorization and SVD have been proposed by Stewart [23], Baker, Gallivan, and Van Dooren [2] and many others, as surveyed in [3]; these ideas are also closely related to rank-revealing QR factorizations [15]. Algorithm 2 differs from those of Stewart in its use of internal pivoting and threshold truncation in place of Stewart's column pivoting. This distinction enables a one-pass algorithm that is closely related to [2, Algorithm 1].

The proposed method is presented in Algorithm 2, which proceeds at each step by orthogonalizing a column of \mathbf{A} against the previously orthogonalized columns. The rank of the resulting factors is controlled through an update-and-delete procedure that is illustrated in Figure 2. After orthogonalizing a column of \mathbf{A} , the algorithm checks if any row of \mathbf{R} has small relative norm; if such a row exists, the corresponding column of \mathbf{Q} makes little contribution to the factorization, so that column of \mathbf{Q} and row of \mathbf{R} can be deleted at only a small loss of accuracy in the factorization. (Future columns of \mathbf{A} will not be orthogonalized against the vector deleted from \mathbf{Q} , so this direction can re-emerge if a later column in \mathbf{A} warrants it.)

Robust implementations of Algorithm 2 should replace the classical Gram–Schmidt operations

$$\mathbf{r} = \mathbf{Q}^T \mathbf{a}, \quad \mathbf{f} = \mathbf{a} - \mathbf{Q}\mathbf{r}$$

with a re-orthogonalization step, as suggested by Daniel, Gragg, Kaufman, and Stewart [8]:

$$\begin{aligned} \mathbf{r} &= \mathbf{Q}^T \mathbf{a} \\ \mathbf{f} &= \mathbf{a} - \mathbf{Q}\mathbf{r} \\ \mathbf{c} &= \mathbf{Q}^T \mathbf{f} \end{aligned} \tag{5.2}$$

$$\mathbf{f} = \mathbf{f} - \mathbf{Q}\mathbf{c} \tag{5.3}$$

$$\mathbf{r} = \mathbf{r} + \mathbf{c} \tag{5.4}$$

$$\rho = \|\mathbf{f}\|$$

$$\mathbf{q} = \mathbf{f}/\rho.$$

The extra steps (5.2)–(5.4) generally provide a \mathbf{Q} that is numerically orthogonal to working precision. Pathological cases are easily overcome with some additional slight modifications; see [13] for a complete analysis. Because this algorithm uses the classical Gram–Schmidt method, one can easily block it for parallel efficiency.

5.1 Incremental QR Error Bounds

At step j the truncation criterion in Algorithm 2 will delete row $\mathbf{r}_i^T = \mathbf{e}_i^T \mathbf{R}_j$ if

$$\|\mathbf{r}_i\| \leq \text{tol} \|\widehat{\mathbf{R}}_j\|_F,$$

where \mathbf{r}_i^T is the row of minimum norm and $\widehat{\mathbf{R}}_j$ denotes \mathbf{R}_j with the i th row deleted. This strategy has a straightforward error analysis, which, in light of the approximation (5.1), also implies an error bound on the resulting SVD.

Lemma 5.1 *Let \mathbf{R}_j be the triangular factor at step j of Algorithm 2, and \mathbf{Q}_j the corresponding orthonormal columns in the approximate QR factorization $\mathbf{A}_j \approx \mathbf{Q}_j \mathbf{R}_j$, where \mathbf{A}_j consists of the first j columns of \mathbf{A} . Then*

$$\|\mathbf{A}_j - \mathbf{Q}_j \mathbf{R}_j\|_F \leq \text{tol} \cdot d_j \cdot \|\mathbf{R}_j\|_F,$$

where d_j is the number of column/row deletions that have been made up to and including step j . (Note that $\mathbf{Q}_j \in \mathbb{R}^{m \times (j-d_j)}$, $\mathbf{R}_j \in \mathbb{R}^{(j-d_j) \times j}$, and $d_n = \min\{m, n\} - k$, where $k = \text{rank}(\mathbf{Q}_n \mathbf{R}_n)$.)

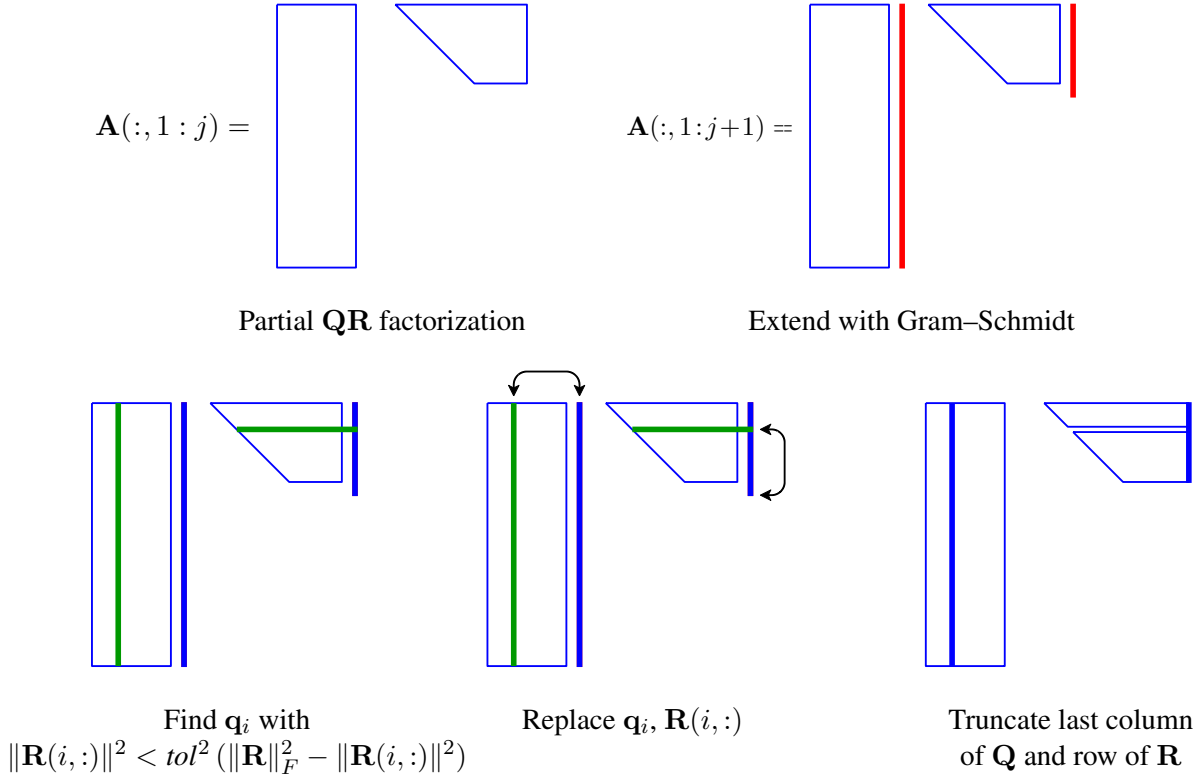


Figure 2: Diagram illustrating the QR update procedure.

Before proving this lemma, we note that it gives a bound on the error in the resulting approximate SVD of \mathbf{A} . Suppose d_n deletions are made when this algorithm computes the approximate factorization $\mathbf{A} \approx \mathbf{Q}\mathbf{R}$ with tolerance tol . Given the SVD $\mathbf{R} = \hat{\mathbf{V}}\mathbf{S}\mathbf{W}^*$, set $\mathbf{V} \equiv \mathbf{Q}\hat{\mathbf{V}}$. Then

$$\|\mathbf{A} - \mathbf{V}\mathbf{S}\mathbf{W}^*\|_F \leq tol \cdot d_n \cdot \|\mathbf{R}\|_F.$$

Proof of Lemma 5.1: The proof shall be by induction. Let $\mathbf{E}_j = \mathbf{A}_j - \mathbf{Q}_j\mathbf{R}_j$ and assume

$$\|\mathbf{E}_j\|_F \leq tol \cdot d_j \cdot \|\mathbf{R}_j\|_F. \quad (5.5)$$

Orthogonalize column $j+1$ of \mathbf{A} using Gram-Schmidt to obtain

$$\mathbf{A}_{j+1} = \mathbf{Q}_{j+1}\mathbf{R}_{j+1} + [\mathbf{E}_j, \mathbf{0}].$$

If no deflation occurs at this step, the bound holds trivially since

$$\|\mathbf{E}_{j+1}\|_F = \|[\mathbf{E}_j, \mathbf{0}]\|_F \leq tol \cdot d_j \cdot \|\mathbf{R}_j\|_F \leq tol \cdot d_{j+1} \cdot \|\mathbf{R}_{j+1}\|_F,$$

because $d_{j+1} = d_j$ and $\|\mathbf{R}_j\|_F \leq \|\mathbf{R}_{j+1}\|_F$.

Suppose \mathbf{R}_j has dimension $k \times j$ (i.e., $k = j - d_j$). Let i be the index of the row of minimum norm and let $\widehat{\mathbf{R}}_{j+1}$ be obtained by deleting the i th row of \mathbf{R}_{j+1} . If $\mathbf{r}_i^T = \mathbf{e}_i^T \mathbf{R}_{j+1}$ satisfies $\|\mathbf{r}_i^T\| \leq \text{tol} \cdot \|\widehat{\mathbf{R}}_{j+1}\|_F$ then deflation occurs. Deleting column i of \mathbf{Q}_{j+1} and row i of \mathbf{R}_{j+1} replaces \mathbf{Q}_{j+1} and \mathbf{R}_{j+1} with $\widehat{\mathbf{Q}}_{j+1}$ and $\widehat{\mathbf{R}}_{j+1}$. Then

$$\widehat{\mathbf{Q}}_{j+1} \widehat{\mathbf{R}}_{j+1} = \mathbf{Q}_{j+1} (\mathbf{R}_{j+1} - \mathbf{e}_i \mathbf{r}_i^T),$$

and

$$\mathbf{A}_{j+1} = \mathbf{Q}_{j+1} (\mathbf{R}_{j+1} - \mathbf{e}_i \mathbf{r}_i^T) + [\mathbf{E}_j, \mathbf{0}] + \mathbf{Q}_{j+1} \mathbf{e}_i \mathbf{r}_i^T.$$

Hence the deletion gives the overall error

$$\mathbf{E}_{j+1} = \mathbf{A}_{j+1} - \widehat{\mathbf{Q}}_{j+1} \widehat{\mathbf{R}}_{j+1} = [\mathbf{E}_j, \mathbf{0}] + \mathbf{Q}_{j+1} \mathbf{e}_i \mathbf{r}_i^T.$$

Therefore, when $i < k + 1$, the inductive assumption (5.5) implies

$$\|\mathbf{E}_{j+1}\|_F \leq \|\mathbf{E}_j\|_F + \|\mathbf{r}_i^T\| \leq \text{tol} \cdot (d_j \cdot \|\mathbf{R}_j\|_F + \|\widehat{\mathbf{R}}_{j+1}\|_F) \leq \text{tol} \cdot (d_j + 1) \cdot \|\widehat{\mathbf{R}}_{j+1}\|_F,$$

since $\widehat{\mathbf{R}}_{j+1}$ contains row $k + 1$ of \mathbf{R}_{j+1} , which must have a norm larger than the row marked for deletion. Since row $k + 1$ of $\widehat{\mathbf{R}}_{j+1}$ consists of just one nonzero element,

$$\|\widehat{\mathbf{R}}_{j+1}\|_F^2 \geq \|\widehat{\mathbf{R}}_j\|_F^2 + \rho_{k+1,j+1}^2 \geq \|\mathbf{R}_j\|_F^2,$$

where $\rho_{k+1,j+1}$ is the element $\mathbf{R}_{j+1}(k+1, j+1)$ and $\widehat{\mathbf{R}}_j$ is the matrix \mathbf{R}_j with i th row deleted. If $i = k + 1$, then the last row of \mathbf{R}_{j+1} is deleted and the desired inequality must hold, since \mathbf{R}_j is a submatrix of $\widehat{\mathbf{R}}_{j+1}$. At the end of this process, replace \mathbf{R}_{j+1} and \mathbf{Q}_{j+1} with $\widehat{\mathbf{R}}_{j+1}$ and $\widehat{\mathbf{Q}}_{j+1}$ to obtain the approximation

$$\|\mathbf{A}_{j+1}\| \leq \text{tol} \cdot d_{j+1} \cdot \|\mathbf{R}_{j+1}\|_F,$$

since $d_{j+1} = d_j + 1$.

The error bound for the base case $j = 1$ clearly holds, completing the induction. ■

The approximate QR factorization that results from this algorithm could be used directly for the approximation of leverage scores. The perturbation theory of Ipsen and Wentworth [18] describes how the tolerance in our algorithm will affect the accuracy of the resulting leverage scores. We also note that for extra expediency this one-pass QR algorithm could be stopped when $\|\widehat{\mathbf{R}}_j\|_F \approx \|\mathbf{A}\|_F$ (at the cost of an extra pass through \mathbf{A} to compute $\|\mathbf{A}\|_F$), or applied to only a random sampling of k columns of \mathbf{A} . (Drmač and Gugercin propose a different random approach to DEIM index selection, based on sampling rows of \mathbf{V} to compute DEIM indices [11].)

6 Computational Examples

This section presents some computational evidence illustrating the excellent approximation properties of the DEIM-CUR factorization, consistent with the error analysis in Section 4. For each of our three examples, we compare the accuracy of the DEIM-CUR factorization with several schemes based on leverage scores. To remove random variations from our experiments, in most cases we select columns and rows having the highest leverage scores; for the first example, we include results for random leverage score sampling. For Example 1 we also study the effect of inaccurate singular vectors on the DEIM selection, and compare the accuracy of DEIM-CUR to CUR approximations based on the column-pivoted QR algorithm.

Example 1. *Low-rank approximation of a sparse, nonnegative matrix*

The first example builds a matrix $\mathbf{A} \in \mathbb{R}^{300,000 \times 300}$ of the form

$$\mathbf{A} = \sum_{j=1}^{10} \frac{2}{j} \mathbf{x}_j \mathbf{y}_j^T + \sum_{j=11}^{300} \frac{1}{j} \mathbf{x}_j \mathbf{y}_j^T, \quad (6.1)$$

where $\mathbf{x}_j \in \mathbb{R}^{300,000}$ and $\mathbf{y}_j \in \mathbb{R}^{300}$ are sparse vectors with random nonnegative entries (in MATLAB, $\mathbf{x}_j = \text{sprand}(300000, 1, 0.025)$ and $\mathbf{y}_j = \text{sprand}(300, 1, 0.025)$). In this instantiation, \mathbf{A} has 15,971,584 nonzeros, i.e., about 18% of all entries are nonzero. The form (6.1) is not a singular value decomposition, since $\{\mathbf{x}_j\}$ and $\{\mathbf{y}_j\}$ are not orthonormal sets; however, this decomposition suggests the structure of the SVD: the singular values decay like $1/j$, and with the first ten singular values weighted more heavily to give a notable drop between σ_{10} and σ_{11} . We begin these experiments by computing \mathbf{V} and \mathbf{W} using MATLAB's economy-sized SVD routine ($[\mathbf{V}, \mathbf{S}, \mathbf{W}] = \text{svd}(\mathbf{A}, '0')$).

Figure 3 compares the error $\|\mathbf{A} - \mathbf{C}\mathbf{U}\mathbf{R}\|$ for DEIM-CUR and methods that take \mathbf{C} and \mathbf{R} as the columns and rows of \mathbf{A} with the highest leverage scores. These scores are computed using either all right and left singular vectors (300 of each), or using only the leading ten right and left singular vectors. Both approaches perform rather worse than DEIM-CUR, which closely tracks the optimal value σ_{k+1} .

To gain insight into these results, we examine the interpolation constants η_p and η_q for all three approaches. Figure 4 shows that these constants are largest for leverage scores based on all the singular vectors; using only ten singular vectors improves both the interpolation constants and the accuracy of the approximation (as seen in Figure 3). The DEIM-CUR method gives better interpolation constants and more accurate approximations.




Figure 3: Accuracy of CUR approximations for the sparse, nonnegative matrix (6.1) using k columns and rows, constructed by DEIM-CUR and two leverage score strategies: “LS (all)” selects rows and columns with highest leverage scores computed using all 300 singular vectors; “LS (10)” only uses the leading ten singular vectors. The “DEIM($\hat{\mathbf{V}}, \hat{\mathbf{W}}$)” curve (nearly atop the “DEIM” curve) uses approximate singular vectors, described later.




Figure 4: Error constants $\eta_p = \|(\mathbf{P}_k^T \mathbf{V}_k)^{-1}\|$ and $\eta_q = \|(\mathbf{W}_k^T \mathbf{Q})^{-1}\|$ for rows and columns selected using two leverage score strategies (left plot) and the DEIM algorithm (right plot), for the matrix (6.1).




Figure 5: Accuracy of CUR approximations for (6.1) generated by randomly sampling rows and columns with probability weighted by leverage scores computed from the leading ten singular vectors. All ten trials (gray lines) perform similarly to the deterministic “LS (10)” approach, and worse than the DEIM-CUR approximation.

A CUR factorization can also be obtained by randomly sampling columns and rows of \mathbf{A} , with the probability of selection weighted by leverage scores [22]. We apply this approach on the current example, selecting $k = 30$ rows and columns of \mathbf{A} with a probability given by the leverage scores computed from the leading ten singular vectors (normalized to give a probability distribution). Figure 5 gives the results of ten independent experiments, showing that while sampling can sometimes yield better results than the deterministic leverage score approach, overall the approximations are still inferior to those from DEIM-CUR.

How robust is the DEIM-CUR approximation to errors in the singular vectors? To investigate, we compute

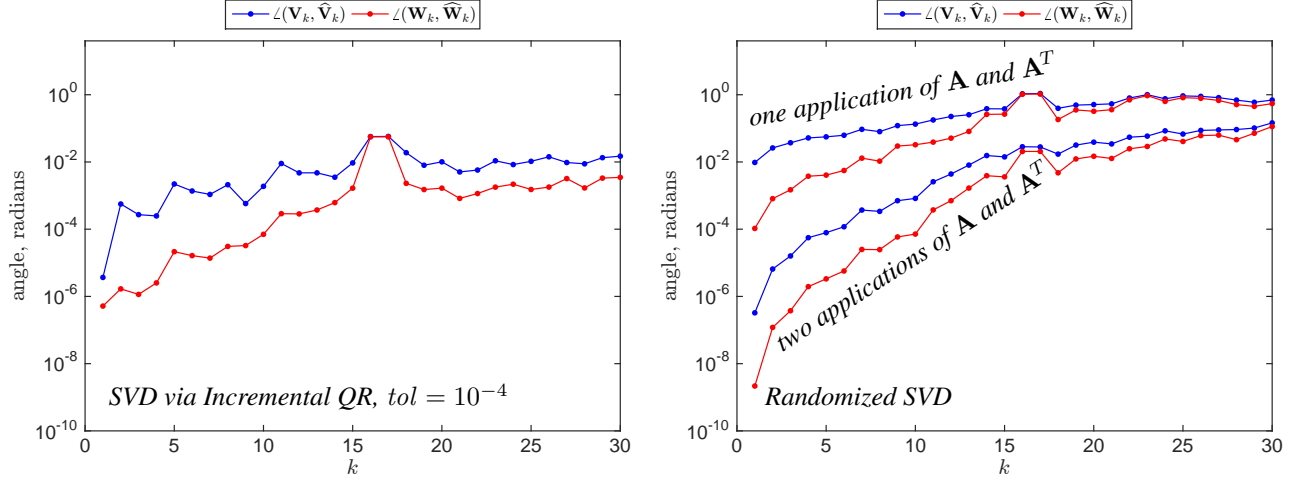


Figure 6: The angle between the leading k -dimensional exact singular subspaces $\text{Ran}(\mathbf{V}_k)$ and $\text{Ran}(\mathbf{W}_k)$ (generated by MATLAB's `svd` command) and approximate singular subspaces $\text{Ran}(\widehat{\mathbf{V}}_k)$ and $\text{Ran}(\widehat{\mathbf{W}}_k)$ for the matrix (6.1). On the left, $\widehat{\mathbf{V}}_k$ and $\widehat{\mathbf{W}}_k$ are generated using the Incremental QR algorithm described in Section 5, with $\text{tol} = 10^{-4}$; on the right, $\widehat{\mathbf{V}}_k$ and $\widehat{\mathbf{W}}_k$ are generated using randomized SVD algorithm [16] using one and two applications of \mathbf{A} and \mathbf{A}^T .

$\widehat{\mathbf{V}} \approx \mathbf{V}$ and $\widehat{\mathbf{W}} \approx \mathbf{W}$ using the Incremental QR algorithm detailed in Section 5 (with $\text{tol} = 10^{-4}$) and the Randomized SVD algorithm described by Halko, Martinsson, and Tropp [16, p. 227]. To give extreme examples of the latter, we compute $\widehat{\mathbf{V}}$ and $\widehat{\mathbf{W}}$ through *only one or two* applications each of \mathbf{A} and \mathbf{A}^T .⁶ As Figure 6 illustrates, in both cases the angle between the exact and approximate leading singular subspaces is significant, particularly as k grows. This drift in the subspaces has little effect on the accuracy of the DEIM approximations.

- The DEIM approximation using the Incremental QR algorithm is quite robust, choosing at most 3 different row indices and 2 different column indices for $k = 1, \dots, 30$, with a relative discrepancy in $\|\mathbf{A} - \mathbf{C}_k \mathbf{U}_k \mathbf{R}_k\|$ of at most 9.27% (and this realized only at step $k = 30$).
- When \mathbf{A} and \mathbf{A}^T are applied once in the Randomized SVD algorithm, the DEIM indices differ considerably from those drawn from exact singular vectors (e.g., for $k = 30$, 20 of 30 row indices and 3 of 30 column indices differ), yet the quality of the approximation $\|\mathbf{A} - \mathbf{C}_k \mathbf{U}_k \mathbf{R}_k\|$ remains almost the same (relative difference of at most 10.45%); see the dashed line in Figure 3.
- When \mathbf{A} and \mathbf{A}^T are applied twice, the DEIM indices are nearly identical (e.g., for $k = 30$, 0 of 30 row indices and 2 of 30 column indices differ). On the scale of the plot in Figure 3, $\|\mathbf{A} - \mathbf{C}_k \mathbf{U}_k \mathbf{R}_k\|$ could not be distinguished from the DEIM-CUR errors using exact singular vectors; the maximum relative discrepancy is 2.21%.

⁶ This corresponds to $q = 0$ and $q = 1$ in the notation of [16, p. 227]. Let the columns of $\mathbf{Q} \in \mathbb{R}^{m \times 2k_{\max}}$ form an orthonormal basis for $(\mathbf{A}\mathbf{A}^T)^q \mathbf{A}\mathbf{\Omega}$, where $\mathbf{\Omega} \in \mathbb{R}^{n \times 2k_{\max}}$ is a random matrix with i.i.d. Gaussian entries and we take $k_{\max} = 30$. Then the leading k_{\max} columns of \mathbf{V} and \mathbf{W} are approximated by taking the SVD of $\mathbf{Q}^* \mathbf{A} \in \mathbb{R}^{2k_{\max} \times n}$.

Thus far we have only compared the DEIM-CUR approximations to CUR factorizations obtained from leverage scores, which also use singular vector information, thus illustrating how DEIM can use the same raw materials to better effect. Next we compare DEIM-CUR to approximations computed using a different approach based on QR factorization of \mathbf{A} ; see, e.g., [7, 23]. Begin by computing a column-pivoted QR factorization of \mathbf{A} ; the first k selected columns give the indices \mathbf{q} , from which we extract $\mathbf{C}_k = \mathbf{A}(:, \mathbf{q})$. Next, a column-pivoted QR




Figure 7: Comparison of DEIM-CUR and QR-CUR performance for 100 sparse random $300,000 \times 300$ matrices of the form (6.1). The top plot shows a histogram of the ratio of $\|\mathbf{A} - \mathbf{C}_k \mathbf{U}_k \mathbf{R}_k\|$ for DEIM-CUR and QR-CUR. The bottom plots compare the error constant $\eta_p = \|(\mathbf{P}_k^T \mathbf{V}_k)^{-1}\|$ for DEIM-CUR (left) and QR-CUR (right); note the logarithmic scale of the horizontal axes in the lower plots.

factorization of \mathbf{C}_k^T is performed; the first k selected columns of \mathbf{C}_k^T give the indices \mathbf{p} , from which we build $\mathbf{R}_k = \mathbf{A}(\mathbf{p}, :)$. We refer to this technique as “QR-CUR.”

Figure 7 compares the results for 100 trials involving sparse random matrices of dimension $300,000 \times 300$ having the form of our first experiment (6.1). DEIM-CUR and QR-CUR produce factorizations with similar accuracy, which we illustrate with a histogram of the ratio of $\|\mathbf{A} - \mathbf{C}_k \mathbf{U}_k \mathbf{R}_k\|$ for DEIM-CUR to QR-CUR, for $k = 1, \dots, 100$. (DEIM-CUR produces a smaller error when the ratio is less than one.) While these errors are similar, the error constants η_p and η_q for the two methods are quite different. The bottom plots in Figure 7 compare histograms of $\log_{10} \eta_p$. For DEIM-CUR, the η_p values are quite consistent across the 100 random \mathbf{A} , while for QR-CUR the η_p values are both larger and rather less consistent. (The figures for η_q are qualitatively identical, but about an order of magnitude smaller for both methods.)

The advantage of DEIM-CUR over approximations based on leverage scores remains when the singular values decrease more sharply. Modify (6.1) to give a more significant drop between σ_{10} and σ_{11} :

$$\mathbf{A} = \sum_{j=1}^{10} \frac{1000}{j} \mathbf{x}_j \mathbf{y}_j^T + \sum_{j=11}^{300} \frac{1}{j} \mathbf{x}_j \mathbf{y}_j^T. \quad (6.2)$$

As seen in Figures 8 and 9, the DEIM-CUR approach again delivers excellent approximations, while selecting the rows and columns with highest leverage scores does not perform nearly as well. (In Figure 9, note the significant jump in the “LS (10)” error constant η_q corresponding to those k values where $\|\mathbf{A} - \mathbf{C}_k \mathbf{U}_k \mathbf{R}_k\|/\sigma_{k+1}$ is large.)




Figure 8: Accuracy of CUR approximations using k rows and columns, for DEIM-CUR and two leverage score strategies for the sparse, nonnegative matrix (6.2). “LS (all)” selects rows and columns having the highest leverage scores computed using all 300 singular vectors; “LS (10)” uses the leading 10 singular vectors.




Figure 9: Error constants $\eta_p = \|(\mathbf{P}_k^T \mathbf{V}_k)^{-1}\|$ and $\eta_q = \|(\mathbf{W}_k^T \mathbf{Q}_k)^{-1}\|$ for rows and columns selected using two leverage score strategies (left plot) and the DEIM algorithm (right plot), for the sparse matrix \mathbf{A} given in (6.2).

Example 2. TechTC term document data

The second example, adapted from Mahoney and Drineas [22], computes the CUR factorization of a term document matrix with data drawn from the Technion Repository of Text Categorization Datasets (TechTC) [12]. The rows of the data matrix correspond to websites (consolidated from multiple webpages), while the columns correspond to “features” (words from the text of the webpages). The (j, k) entry of \mathbf{A} reflects the importance of the feature text on the given website; most entries are zero. For this experiment we use TechTC-100 test set 26, which concatenates a data set relating to Evansville, Indiana (id 10567) with another for Miami, Florida (id 11346). Following Mahoney and Drineas [22], we omit all features with four or fewer characters from the data set, leaving a matrix with 139 rows and 15,170 columns. Each row of \mathbf{A} is then scaled to have unit 2-norm. Ideally a CUR factorization not only gives an accurate low-rank approximation to \mathbf{A} , but also selects rows corresponding to representative webpages from each geographic area, and columns corresponding to meaningful features.

Figure 10 compares DEIM-CUR approximations to row and column selection based on highest leverage scores (from all singular vectors, or the two leading singular vectors). The DEIM-CUR approximations are typically more accurate than those based on leverage scores, but all approaches give errors roughly two times larger than the slowly-decaying optimal value of σ_{k+1} . How do the DEIM columns (features) compare to those with the highest leverage scores? Figure 11 shows the leverage scores associated with each column of \mathbf{A} (based on the two leading singular vectors), along with the first 30 columns selected by DEIM. While the columns with highest leverage scores were found by DEIM, there are DEIM columns with marginal leverage scores, and vice versa. This data is more easily parsed in Table 1, which lists the features corresponding to the first 20 DEIM columns. (To ease comparison, we normalize leverage scores so that the maximum value is one.) The leading features identified by DEIM, including “evansville” (first DEIM point), “florida” (second), “miami” (sixth), and “indiana” (nineteenth), indeed reveal key geographic terms. These terms scored at least as high when ranked by leverage scores based on two leading singular vectors; when all singular vectors are used, the scores of these terms generally drop, relative to other features. Overall, one notes that DEIM selects a significantly different set of indices than those valued by leverage scores, and, as seen in Figure 10, tends to provide a somewhat better low-rank approximation.




Figure 10: Accuracy of CUR factorizations for the TechTC example, selecting rows and columns using top leverage scores for all singular vectors and the leading two singular vectors, and DEIM.




Figure 11: The columns selected by DEIM for the TechTC example, compared to leverage scores from the leading two singular vectors.

Example 3. Tumor detection in genetics data

Our final example uses the GSE10072 cancer genetics data set from the National Institutes of Health, previously investigated by Kundu, Nambirijan, and Drineas [19]. The matrix $\mathbf{A} \in \mathbb{R}^{22,283 \times 107}$ contains data for 22,283 probes applied to 107 patients. The (j, k) entry of \mathbf{A} reflects how strongly patient k responded to probe j . This experiment seeks probes that segment the population into two clusters: the 58 patients with tumors, and the 49 without.⁷ To

⁷The data is available from <http://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE10072>.

Table 1: The features selected by DEIM-CUR for the TechTC data set, compared to the (scaled) leverage scores using the leading two singular vectors, and all singular vectors.

DEIM rank	index q_j	LS (2)		LS (all)		feature
		rank	score	rank	score	
1	10973	1	1.000	4	0.875	evansville
2	1	2	0.741	8	0.726	florida
3	1547	13	0.031	2	0.948	spacer
4	109	8	0.055	66	0.347	contact
5	209	12	0.040	32	0.458	service
6	50	4	0.116	6	0.739	miami
7	824	46	0.007	5	0.809	chapter
8	1841	33	0.010	20	0.537	health
9	171	5	0.113	13	0.617	information
10	234	16	0.026	37	0.436	events
11	595	84	0.004	15	0.576	church
12	60	15	0.026	67	0.347	email
13	945	10	0.047	30	0.474	services
14	1670	129	0.002	1	1.000	bullet
15	216	35	0.009	38	0.430	music
16	78	3	0.246	24	0.492	south
17	213	19	0.018	110	0.259	their
18	138	14	0.030	43	0.408	please
19	6110	7	0.060	95	0.280	indiana
20	1152	70	0.005	152	0.221	member

center the data, we subtract the mean of each row from all entries in that row. As shown in [19], the leading two principal vectors of this matrix segment the population very well.

Like the TechTC data, the singular values of \mathbf{A} decay slowly, as seen in Figure 12. Once again the DEIM-CUR procedure produces a more accurate low-rank approximation than obtained by selecting the rows and columns with highest leverage scores, whether those are computed using all the singular vectors, or just the leading two or ten.

Table 2 reports the first 15 rows selected by the DEIM-CUR process, along with the corresponding leverage scores based on two, ten, and all singular vectors. Do the probes selected by DEIM discriminate the patients with tumors (“sick”) from those without (“well”)? To investigate, for each selected probe we count the number of large positive entries corresponding to sick and well patients.⁸ Some but not all of the DEIM-CUR probes effectively select only sick or well patients. Contrast these results with Table 3, which shows the probes with highest leverage scores (drawn from the leading two singular vectors). Only four of these probes were also selected by the DEIM procedure (even if we continue the DEIM procedure to select the maximum number, $n = 107$, of indices). This discrepancy is quite different from the good agreement between DEIM and leverage score indices for the TechTC data in Table 1, despite the similar dimensions and the comparably slow decay of the singular values.

While the rows selected from leverage scores did not produce as accurate an approximation, $\|\mathbf{A} - \mathbf{C}_k \mathbf{U}_k \mathbf{R}_k\|$,

⁸In particular, we call an entry of the mean-centered matrix \mathbf{A} large if its value exceeds one. Of the 22,283 probes, for only 23 probes do at least 30 of the 58 sick patients have such large entries; for only 95 probes do at least 30 of the 49 well patients have large entries. There is no overlap between the probes that are strongly expressed by the sick and well patients.




Figure 12: Accuracy of CUR factorizations for a genetics data set. DEIM-CUR consistently outperforms factorizations derived by taking the rows and columns with largest leverage scores, regardless of whether these scores are drawn from all singular vectors, the leading ten singular vectors, or the leading two singular vectors.

Table 2: Genetics example: the probes selected by DEIM-CUR, compared to the (scaled) leverage scores using the leading two singular vectors, ten singular vectors, and all singular vectors.

DEIM rank	index q_j	probe set	gene name	number sick	number well	LS (2)		LS (10)		LS (all)	
						rank	score	rank	score	rank	score
1	9565	210081_at	AGER	2	45	1	1.000	45	0.504	386	0.123
2	14270	214895_s.at	ADAM10	8	3	211	0.173	1171	0.108	3344	0.036
3	8650	209156_s.at	COL6A2	5	6	15156	0.005	252	0.245	708	0.091
4	11057	211653_x.at	AKR1C2	18	1	6440	0.017	11	0.656	146	0.185
5	14153	214777_at	IGKV4-1	27	3	281	0.148	19	0.607	106	0.209
6	18976	219612_s.at	FGG	17	17	2591	0.039	2	0.956	4	0.825
7	3831	204304_s.at	PROM1	16	4	992	0.073	70	0.417	32	0.345
8	3351	203824_at	TSPAN8	17	4	9687	0.011	21	0.582	31	0.355
9	4275	204748_at	PTGS2	18	14	424	0.118	13	0.624	42	0.313
10	1437	201909_at	RPS4Y1	21	34	8232	0.013	3	0.913	5	0.736
11	14150	214774_x.at	TOX3	34	0	95	0.262	49	0.492	102	0.210
12	10518	211074_at	FOLR1	7	4	9482	0.011	926	0.124	213	0.159
13	9580	210096_at	CYP4B1	6	44	8	0.797	65	0.431	54	0.284
14	4002	204475_at	MMP1	27	0	34	0.406	24	0.564	21	0.465
15	13990	214612_x.at	MAGEA	16	0	489	0.110	134	0.323	35	0.339

as DEIM, these probes do a much more effective job of discriminating patients with tumors from those without. Indeed, for 14 of the top 15 probes, the tumor-free patients express strongly, while the patients with tumors do not; in the remaining case, the opposite occurs.

Table 3: Genetics example: the probes with top (scaled) leverage scores, derived from the first two singular vectors.

LS (2) rank	index q_j	LS (2) score	probe set	gene name	number sick	number well	DEIM rank
1	9565	1.000	210081_at	AGER	2	45	1
2	13766	0.922	214387_x_at	SFTPC	6	48	—
3	11135	0.907	211735_x_at	SFTPC	5	48	73
4	9361	0.899	209875_s_at	SPP1	50	2	—
5	5509	0.896	205982_x_at	SFTPC	5	48	—
6	9103	0.835	209613_s_at	ADH1B	2	47	—
7	14827	0.834	215454_x_at	SFTPC	0	46	—
8	9580	0.797	210096_at	CYP4B1	6	44	13
9	4239	0.754	204712_at	WIF1	5	43	70
10	3507	0.724	203980_at	FABP4	2	44	—
11	18594	0.717	219230_at	TMEM100	2	38	—
12	9102	0.684	209612_s_at	ADH1B	2	46	—
13	13514	0.626	214135_at	CLDN18	3	47	—
14	5393	0.626	205866_at	FCN3	0	39	—
15	4727	0.614	205200_at	CLEC3B	0	39	—

7 Conclusions

The Discrete Empirical Interpolation Method (DEIM) is an index selection procedure that gives simple, deterministic CUR factorizations of the matrix \mathbf{A} . Since DEIM utilizes (approximate) singular vectors, we propose an effective one-pass incremental approximate QR factorization that can efficiently compute dominant singular vectors for data sets with rapidly decaying singular values; this method could prove useful in a variety of other settings. The accuracy of the resulting rank- k CUR factorization can be bounded in terms of σ_{k+1} , the error in the best rank- k approximation to \mathbf{A} . Our analysis of the 2-norm error $\|\mathbf{A} - \mathbf{CUR}\|$ applies to all CUR approximations that use the optimal central factor $\mathbf{U} = \mathbf{C}^T \mathbf{A} \mathbf{R}^T$, and hence can give insight into the performance of other index selection algorithms, such as leverage scores, uniform random sampling, or column-pivoted QR factorization. Numerical examples illustrate that the DEIM-CUR approach can deliver very good low rank approximations, compared to row selection based on dominant leverage scores.

Acknowledgements

We thank Inderjit Dhillon, Petros Drineas, Ilse Ipsen, Michael Mahoney, and Nick Trefethen for a number of helpful discussions. We are also grateful to Gunnar Martinsson for recommending experiments with the column-pivoted QR selection algorithm, and to an anonymous referee for encouraging us to seek the improved analysis and growth example for the DEIM error constants at the end of Section 4.

References

- [1] James Baglama and Lothar Reichel. An implicitly restarted block Lanczos bidiagonalization method using Leja shifts. *BIT*, 53:285–310, 2013.
- [2] C. G. Baker, K. A. Gallivan, and P. Van Dooren. Low-rank incremental methods for computing dominant singular subspaces. *Linear Algebra Appl.*, 436:2866–2888, 2012.
- [3] Christopher G. Baker. A block incremental algorithm for computing dominant singular subspaces. Master’s thesis, Florida State University, 2004.
- [4] Maxime Barrault, Yvon Maday, Ngoc Cuong Nguyen, and Anthony T. Patera. An ‘empirical interpolation’ method: application to efficient reduced-basis discretization of partial differential equations. *Comptes Rendus Acad. Sci. Paris, Ser. I*, 339:667–672, 2004.
- [5] Christos Boutsidis and David P. Woodruff. Optimal CUR matrix decompositions. arXiv:1405.7910 [cs.DS], 2014.
- [6] Saifon Chaturantabut and Danny C. Sorensen. Nonlinear model reduction via discrete empirical interpolation. *SIAM J. Sci. Comput.*, 32:2737–2764, 2010.
- [7] H. Cheng, Z. Gimbutas, P. G. Martinsson, and V. Rokhlin. On the compression of low rank matrices. *SIAM J. Sci. Comput.*, 26:1389–1404, 2005.
- [8] J. W. Daniel, W. B. Gragg, L. Kaufman, and G. W. Stewart. Reorthogonalization and stable algorithms for updating the Gram–Schmidt QR factorization. *Math. Comp.*, 30:772–795, 1976.
- [9] Jack J. Dongarra, Iain S. Duff, Danny C. Sorensen, and Henk A. van der Vorst. *Numerical Linear Algebra for High-Performance Computers*. SIAM, Philadelphia, 1998.
- [10] Petros Drineas, Michael W. Mahoney, and S. Muthukrishnan. Relative-error CUR matrix decompositions. *SIAM J. Matrix Anal. Appl.*, pages 844–881, 2008.
- [11] Zlatko Drmač and Serkan Gugercin. A new selection operator for the discrete empirical interpolation method — improved a priori error bound and extensions. arXiv:1505.0037 [cs.NA], 2015.
- [12] Evgeniy Gabrilovich and Shaul Markovitch. Text categorization with many redundant features: Using aggressive feature selection to make SVMs competitive with C4.5. In *The 21st International Conference on Machine Learning (ICML)*, pages 321–328, 2004.
- [13] Luc Giraud, Julien Langou, Miroslav Rozložník, and Jasper van den Eshof. Rounding error analysis of the classical Gram–Schmidt orthogonalization process. *Numer. Math.*, 101:87–100, 2005.
- [14] S. A. Goreinov, E. E. Tyrtyshnikov, and N. L. Zamarashkin. A theory of pseudoskeleton approximations. *Linear Algebra Appl.*, 261:1–21, 1997.

- [15] Ming Gu and Stanley C. Eisenstat. Efficient algorithms for computing a strong rank-revealing QR factorization. *SIAM J. Sci. Comput.*, 17:848–869, 1996.
- [16] N. Halko, P. G. Martinsson, and J. A. Tropp. Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53:217–288, 2011.
- [17] Michiel E. Hochstenbach. A Jacobi–Davidson type SVD method. *SIAM J. Sci. Comput.*, 23:606–628, 2001.
- [18] Ilse C. F. Ipsen and Thomas Wentworth. Sensitivity of leverage scores. arXiv:1402.0957 [math.NA], 2014.
- [19] Abhisek Kundu, Srinivas Nambirajan, and Petros Drineas. Identifying influential entries in a matrix. arXiv:1310.3556 [cs.nA], 2013.
- [20] Rasmus Munk Larsen. PROPACK: Software for large and sparse SVD calculations. <http://sun.stanford.edu/~rmunk/PROPACK/>, 2005. Version 2.1.
- [21] R. B. Lehoucq, D. C. Sorensen, and C. Yang. *ARPACK Users’ Guide: Solution of Large-Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods*. SIAM, Philadelphia, 1998.
- [22] Michael W. Mahoney and Petros Drineas. CUR matrix decompositions for improved data analysis. *Proc. Nat. Acad. Sci.*, 106:697–702, 2009.
- [23] G. W. Stewart. Four algorithms for the efficient computation of truncated QR approximations to a sparse matrix. *Numer. Math.*, 83:313–323, 1999.
- [24] Daniel B. Szyld. The many proofs of an identity on the norm of oblique projections. *Numer. Alg.*, 42:309–323, 2006.
- [25] Christian Thureau, Kristian Kersting, and Christian Bauckhage. Deterministic CUR for improved large-scale data analysis: an empirical study. In *Proceedings of the 2012 SIAM International Conference on Data Mining*, pages 684–695, 2012.
- [26] Lloyd N. Trefethen and David Bau, III. *Numerical Linear Algebra*. SIAM, Philadelphia, 1997.
- [27] Lloyd N. Trefethen and Robert S. Schreiber. Average-case stability of Gaussian elimination. *SIAM J. Matrix Anal. Appl.*, 11:335–360, 1990.
- [28] Shusen Wang and Zhihua Zhang. Improving CUR matrix decomposition and the Nyström approximation via adaptive sampling. *J. Machine Learning Res.*, 14:2729–2769, 2013.